Mathematics of Stochastic Processes & Randomness

Volume I

Markov Chains, Mixing & Convergence

Silvano Salvador

2026

Contents

Preface iv

I Foundations: Probability Spaces and Transition Kernels 1

1 Probability Spaces and Measure-Theoretic Prerequisites 2

1.1 Sigma-Algebras, Measures, and Probability Spaces 2

1.1.1 Random Variables and Distributions 3

1.1.2 Conditional Probability and Expectation 3

1.2 Stochastic Processes: Definitions and Filtrations 3

1.3 Independence and Product Measures 4

2 Markov Chains: Definition and Basic Structure 5

2.1 The Markov Property 5

2.2 Transition Kernels and the Chapman–Kolmogorov Equations 6

2.3 Communication Classes and Irreducibility 6

2.4 Periodicity 7

2.5 First-Passage Decompositions and the Green's Function 7

2.6 Stopping Times and the Strong Markov Property 9

II Long-Run Behavior: Stationarity and Ergodicity 12

3 Stationary Distributions 13

3.1 Definition and Existence 13

3.2 Detailed Balance and Reversibility 14

3.3 Finding Stationary Distributions 14

3.4 Transience and Recurrence 15

3.5 Variance of Additive Functionals and the Asymptotic Variance Formula 15

4 Ergodic Theory for Markov Chains 18

4.1 The Ergodic Theorem 18

4.2 Convergence to Stationarity 18

III Spectral Theory and the Spectral Gap 20

5 Spectral Analysis of Transition Matrices 21
5.1 Eigenvalues of Stochastic Matrices 21
5.2 The Spectral Gap 21
5.3 Canonical Paths and the Congestion Ratio 22
5.4 Cheeger's Inequality 24

IV Coupling Methods and Mixing Time Bounds 25

6 The Coupling Technique 26
6.1 Definition and Basic Properties 26
6.2 Markovian Couplings 26
6.3 The Mixing Time 27
6.4 Doeblin's Coupling and Quantitative Convergence 27

7 Advanced Coupling Techniques 30
7.1 Monotone Coupling and Stochastic Ordering 30
7.2 Coupling from the Past 30

V Continuous-Time Markov Chains 32

8 Generators and Semigroups 33
8.1 Definition of Continuous-Time Markov Chains 33
8.2 The Generator 33
8.3 Stationary Distribution and Reversibility 34

9 Reversible Chains and Spectral Methods in Continuous Time 35
9.1 The Relaxation Time 35
9.2 Log-Sobolev Inequalities 35

VI Applications: MCMC, Random Walks, and Beyond 37

10 Markov Chain Monte Carlo 38
10.1 The Metropolis–Hastings Algorithm 38
10.2 Gibbs Sampling 39

10.3 Mixing Times for MCMC 39

11 Random Walks on Graphs **40**
11.1 The Lazy Random Walk 40
11.2 Spectral Gap for Random Walks 40
11.3 Expander Graphs 41

12 Birth-Death Chains and Queueing Theory **42**
12.1 Spectral Theory of Birth-Death Chains 42
12.2 The M/M/1 Queue 42
12.3 Hitting Times and Fundamental Matrix 43

13 Convergence Rates: Lower Bounds and Cutoff Phenomenon **44**
13.1 Lower Bounds on Mixing Times 44
13.2 The Cutoff Phenomenon 45
13.3 Product Chains and Tensorization 45

14 Martingales and Optional Stopping **47**
14.1 Martingales 47
14.2 Azuma–Hoeffding and McDiarmid Inequalities 48
14.3 Central Limit Theorem for Markov Chains 48

15 Markov Chains in Statistical Physics **49**
15.1 The Ising Model 49
15.2 The Swendsen–Wang and Heat-Bath Algorithms 49

16 Harris Chains and General State Spaces **51**
16.1 Minorization Conditions 51
16.2 Regeneration Theory 52

17 Computational Algorithms **53**
17.1 Pseudocode for Core Algorithms 53
17.2 Diagnostics and Convergence Assessment 54

18 Spectral Methods: Advanced Topics **55**
18.1 Expander Mixing Lemma 55
18.2 Ramanujan Graphs 55
18.3 Geometric Random Walks 56
18.4 Langevin Monte Carlo 56

19 Stochastic Dominance and Monotone Chains **57**

19.1 Stochastic Ordering on Lattices . . . 57
19.2 Monotone Coupling and Sandwiching . . . 57
19.3 Interacting Particle Systems . . . 58

20 Selected Open Problems and Research Directions **59**
20.1 The KLS Conjecture . . . 59
20.2 Mixing and Computation . . . 59
20.3 Non-Reversible Chains and Acceleration . . . 60

21 Renewal Theory and Regenerative Processes **61**
21.1 Renewal Sequences . . . 61
21.2 The Renewal Theorem . . . 62
21.3 Regenerative Processes and the Renewal-Reward Theorem . . . 62
21.4 Renewal Theory and Markov Chain Hitting Times . . . 63
21.5 Markov Renewal Processes . . . 64

22 Large Deviations for Markov Chains **65**
22.1 The Large Deviation Principle . . . 65
22.2 The Donsker–Varadhan Variational Formula . . . 66
22.3 Sanov's Theorem and Sample Path Large Deviations . . . 67
22.4 Applications: Exit Times and Metastability . . . 67

23 Markov Chains on Infinite and Random Graphs **69**
23.1 Recurrence and Transience Revisited: The Role of Geometry . . . 69
23.2 Random Walk on Random Graphs: The Erdős–Rényi Model . . . 70
23.3 Random Walk on Galton–Watson Trees . . . 70
23.4 The Configuration Model and Local Weak Convergence . . . 71

24 Hidden Markov Models and Filtering **73**
24.1 The Hidden Markov Model Framework . . . 73
24.2 The Forward Algorithm and Filtering Recursion . . . 73
24.3 Smoothing and the Forward–Backward Algorithm . . . 74
24.4 The Viterbi Algorithm: Most Likely State Sequence . . . 75
24.5 Parameter Estimation: The Baum–Welch Algorithm . . . 75
24.6 Particle Filtering for Continuous and Large State Spaces . . . 76

25 Markov Decision Processes and Optimal Control **78**
25.1 The Markov Decision Process Framework . . . 78
25.2 Discounted Reward Criterion and the Bellman Equation . . . 78

25.3 Value Iteration and Policy Iteration 79
25.4 Average-Reward MDPs and the Connection to Stationary Distributions 80
25.5 Reinforcement Learning: Learning Without a Known Model 80
25.6 A Worked Example: The Optimal Stopping Problem 81

Concluding Remarks **84**

A Linear Algebra and Perron–Frobenius Theory **85**
A.1 Eigenvalues and Eigenvectors of Stochastic Matrices 85
A.2 Diagonalization and Spectral Decomposition 86
A.3 Doubly Stochastic Matrices and the Birkhoff–von Neumann Theorem . 86

B Glossary of Notation and Standard Distributions **88**
B.1 Notation Conventions 88
B.2 Standard Discrete Distributions Used in Examples 88

C Pseudocode Reference for Key Algorithms **90**

D Compendium of Mixing Time Bounds and Concentration Inequalities **92**
D.1 Master Table of Mixing Time Bounds 92
D.2 Concentration Inequalities Used Throughout the Volume 93
D.3 Cross-Reference Table: Where Each Technique Is Developed 94

E Selected Exercises **96**
E.1 Foundations and Basic Structure (Chapters 1–3) 96
E.2 Stationarity, Ergodicity, and Spectral Theory (Chapters 3–5) 97
E.3 Coupling and Mixing Time Bounds (Chapters 6–7, 14) 97
E.4 Continuous-Time Chains and MCMC (Chapters 8–10) 98
E.5 Martingales, Statistical Physics, and General State Spaces (Chapters 16–17) 98
E.6 Advanced Topics: Renewal, Large Deviations, Random Graphs, HMMs, and MDPs (Chapters 19–23) 98

Bibliography **101**

Preface

This volume opens the series *Mathematics of Stochastic Processes & Randomness* with a rigorous treatment of Markov chains in discrete and continuous time. The theory of Markov chains stands at the intersection of probability, linear algebra, and functional analysis, and constitutes the conceptual foundation for stochastic calculus, random fields, and modern applications ranging from statistical physics to reinforcement learning.

The presentation is self-contained but mathematically demanding. Each concept is introduced via a precise definition, illustrated with worked examples, and developed into a body of theorems with complete proofs. The reader is expected to be comfortable with real analysis, linear algebra, and elementary measure theory. The progression follows the classical hierarchy: transition kernels, irreducibility and periodicity, stationary distributions, convergence theorems, spectral methods, and the quantitative theory of mixing times.

Part I establishes the measure-theoretic foundations and the algebraic structure of transition operators. Part II develops the long-run theory of ergodicity and stationary measures. Part III introduces the spectral gap and its fundamental role in bounding convergence rates. Part IV presents the coupling technique, one of the most powerful probabilistic tools, and derives sharp mixing-time bounds. Part V extends the theory to continuous-time chains and reversible processes. Part VI surveys modern applications including MCMC algorithms, random walks on graphs, and birth-death processes.

Part I

Foundations: Probability Spaces and Transition Kernels

Chapter 1

Probability Spaces and Measure-Theoretic Prerequisites

1.1 Sigma-Algebras, Measures, and Probability Spaces

The mathematical framework underpinning all of probability theory is the *probability space*, a triple $(\Omega, \mathcal{F}, \mathbb{P})$ where every statement about randomness is made precise.

Definition: Probability Space

A **probability space** is a triple $(\Omega, \mathcal{F}, \mathbb{P})$ where:

(i) Ω is a non-empty set called the **sample space**;

(ii) $\mathcal{F}$ is a σ-algebra on Ω, i.e. a collection of subsets closed under complementation and countable unions, with $\emptyset, \Omega \in \mathcal{F}$;

(iii) $\mathbb{P} : \mathcal{F} \to [0, 1]$ is a **probability measure**: σ-additive with $\mathbb{P}(\Omega) = 1$.

The requirement of σ-additivity is the core axiom: for any countable collection of pairwise disjoint events $\{A_n\}_{n \geq 1} \subset \mathcal{F}$,

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mathbb{P}(A_n).$$

This axiom captures the principle that probability is additive over independent contributions and extends naturally to limiting arguments.

Example: Finite State Space

Let $\Omega = \{1, 2, \ldots, N\}$, $\mathcal{F} = 2^{\Omega}$ (all subsets), and $\mathbb{P}(\{i\}) = p_i \geq 0$ with $\sum_{i=1}^{N} p_i = 1$. This is the setting for *finite Markov chains*, where the state space is explicit and all probabilities are encoded in a vector $\boldsymbol{p} = (p_1, \ldots, p_N)^{\top}$.

1.1.1 Random Variables and Distributions

A **random variable** is a measurable function $X : (\Omega, \mathcal{F}) \to (\mathcal{X}, \mathcal{G})$, where $(\mathcal{X}, \mathcal{G})$ is a measurable space. The **distribution** (or law) of X is the pushforward measure $\mu_X = \mathbb{P} \circ X^{-1}$, defined by $\mu_X(B) = \mathbb{P}(X \in B)$ for all $B \in \mathcal{G}$.

When $\mathcal{X}$ is countable (finite or countably infinite), X is called a **discrete random variable** and its distribution is described by the probability mass function $p(x) = \mathbb{P}(X = x)$. When $\mathcal{X} = \mathbb{R}$ and μ_X admits a density with respect to Lebesgue measure, X is **absolutely continuous** with density f_X satisfying $\mathbb{P}(X \in B) = \int_B f_X(x)\,dx$.

1.1.2 Conditional Probability and Expectation

For $A, B \in \mathcal{F}$ with $\mathbb{P}(B) > 0$, the **conditional probability** of A given B is

$$\mathbb{P}(A \mid B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}.$$

The **expected value** of an integrable random variable X is the Lebesgue integral $\mathbb{E}[X] = \int_\Omega X(\omega)\,d\mathbb{P}(\omega)$. **Conditional expectation** $\mathbb{E}[X \mid \mathcal{G}]$ with respect to a sub-σ-algebra $\mathcal{G} \subset \mathcal{F}$ is the unique (up to $\mathbb{P}$-a.s. equivalence) $\mathcal{G}$-measurable random variable satisfying

$$\int_G \mathbb{E}[X \mid \mathcal{G}]\,d\mathbb{P} = \int_G X\,d\mathbb{P} \qquad \forall G \in \mathcal{G}.$$

This is the Radon–Nikodym derivative of the measure $A \mapsto \mathbb{E}[X \cdot \mathbf{1}_A]$ with respect to $\mathbb{P}|_{\mathcal{G}}$.

1.2 Stochastic Processes: Definitions and Filtrations

Definition: Stochastic Process

A **stochastic process** indexed by $T \subseteq [0, \infty)$ or $T = \mathbb{N}_0$ is a family $\{X_t\}_{t \in T}$ of random variables defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$ taking values in a measurable state space $(\mathcal{X}, \mathcal{G})$.

The **natural filtration** generated by $\{X_t\}$ is $\mathcal{F}_t = \sigma(X_s : s \le t)$, the smallest σ-algebra making all X_s for $s \le t$ measurable. A process $\{X_t\}$ is **adapted** to a filtration $\{\mathcal{F}_t\}$ if X_t is $\mathcal{F}_t$-measurable for each t.

The **path** (or **trajectory**) of a stochastic process is the function $t \mapsto X_t(\omega)$ for a fixed $\omega \in \Omega$. The collection of all paths forms the **path space**, which is the natural domain for studying sample-path properties such as continuity, right-continuity, and the existence of limits from the left.

Remark: Kolmogorov Extension Theorem

The Kolmogorov extension theorem guarantees that any consistent family of finite-dimensional distributions $\{\mu_{t_1,\dots,t_n}\}$ uniquely determines a stochastic pro-

cess on the product space $\mathcal{X}^T$. This foundational result justifies defining Markov chains by specifying only their transition probabilities without constructing the underlying probability space explicitly.

1.3 Independence and Product Measures

Events $A_1, \ldots, A_n \in \mathcal{F}$ are **mutually independent** if

$$\mathbb{P}\left(\bigcap_{i \in S} A_i\right) = \prod_{i \in S} \mathbb{P}(A_i)$$

for every non-empty $S \subseteq \{1, \ldots, n\}$. Random variables $X_1, \ldots, X_n$ are independent if the σ-algebras $\sigma(X_1), \ldots, \sigma(X_n)$ are mutually independent.

The **product measure** $\mu_1 \otimes \cdots \otimes \mu_n$ on $(\mathcal{X}_1 \times \cdots \times \mathcal{X}_n, \mathcal{G}_1 \otimes \cdots \otimes \mathcal{G}_n)$ is the unique measure satisfying $(\mu_1 \otimes \cdots \otimes \mu_n)(B_1 \times \cdots \times B_n) = \prod_i \mu_i(B_i)$. The joint distribution of independent random variables $X_i \sim \mu_i$ is exactly the product measure.

Chapter 2

Markov Chains: Definition and Basic Structure

2.1 The Markov Property

The Markov property formalizes the notion that the future of a stochastic process, given its present state, is independent of its past. This memorylessness is the defining characteristic of Markov chains.

Definition: Discrete-Time Markov Chain

A sequence $\{X_n\}_{n\geq 0}$ of $\mathcal{X}$-valued random variables on $(\Omega, \mathcal{F}, \mathbb{P})$ is a **Markov chain** with state space $\mathcal{X}$ if, for all $n \geq 0$ and all measurable $A \subseteq \mathcal{X}$:

$$\mathbb{P}(X_{n+1} \in A \mid X_0, X_1, \ldots, X_n) = \mathbb{P}(X_{n+1} \in A \mid X_n) \quad \mathbb{P}\text{-a.s.}$$

The chain is **time-homogeneous** if $\mathbb{P}(X_{n+1} \in A \mid X_n = x)$ does not depend on n.

For a time-homogeneous chain on a countable state space $\mathcal{X}$, all the information about the dynamics is encoded in the **transition matrix** (or **stochastic matrix**) $P = (p_{ij})_{i,j\in\mathcal{X}}$, where

$$p_{ij} = \mathbb{P}(X_{n+1} = j \mid X_n = i),$$

satisfying $p_{ij} \geq 0$ for all i, j and $\sum_{j\in\mathcal{X}} p_{ij} = 1$ for all i. Such a matrix is called a **row-stochastic matrix**.

Example: Simple Random Walk on $\mathbb{Z}$

Let $\mathcal{X} = \mathbb{Z}$. The **simple symmetric random walk** has transition probabilities $p_{i,i+1} = p_{i,i-1} = 1/2$ for all $i \in \mathbb{Z}$ and $p_{ij} = 0$ otherwise. Starting from $X_0 = 0$, after n steps, $X_n = \sum_{k=1}^{n} \xi_k$ where $\{\xi_k\}$ are i.i.d. with $\mathbb{P}(\xi_k = \pm 1) = 1/2$. By the central limit theorem, $X_n/\sqrt{n} \Rightarrow \mathcal{N}(0,1)$. This process is the prototype for Brownian motion.

2.2 Transition Kernels and the Chapman–Kolmogorov Equations

The n-step transition probabilities $p_{ij}^{(n)} = \mathbb{P}(X_n = j \mid X_0 = i)$ satisfy the fundamental **Chapman–Kolmogorov equations**:

Theorem: Chapman–Kolmogorov

For any $m, n \geq 0$ and all states $i, j \in \mathcal{X}$:

$$p_{ij}^{(m+n)} = \sum_{k \in \mathcal{X}} p_{ik}^{(m)} \, p_{kj}^{(n)}.$$

Equivalently, in matrix form: $P^{m+n} = P^m \cdot P^n$, so $P^{(n)} = P^n$ (the n-th matrix power).

Proof. By the law of total probability and the Markov property applied at time m:

$$\begin{aligned} p_{ij}^{(m+n)} &= \mathbb{P}(X_{m+n} = j \mid X_0 = i) \\ &= \sum_{k \in \mathcal{X}} \mathbb{P}(X_{m+n} = j \mid X_m = k) \, \mathbb{P}(X_m = k \mid X_0 = i) \\ &= \sum_{k \in \mathcal{X}} p_{kj}^{(n)} \, p_{ik}^{(m)}. \end{aligned}$$

The second equality uses the Markov property: given $X_m = k$, the future evolution from time m is independent of the past $X_0, \ldots, X_{m-1}$, and the time-homogeneity ensures $\mathbb{P}(X_{m+n} = j \mid X_m = k) = p_{kj}^{(n)}$. □

2.3 Communication Classes and Irreducibility

Definition: Accessibility and Communication

State j is **accessible** from state i, written $i \to j$, if $p_{ij}^{(n)} > 0$ for some $n \geq 0$. States i and j **communicate**, written $i \leftrightarrow j$, if $i \to j$ and $j \to i$. Communication is an equivalence relation, and its equivalence classes are the **communication classes** of the chain.

A Markov chain is **irreducible** if it has a single communication class — every state is accessible from every other state. Irreducibility is a crucial structural assumption for many convergence results.

Proposition: Irreducibility and Positivity

A finite Markov chain with transition matrix P is irreducible if and only if for every pair (i, j) there exists $n = n(i, j) \geq 1$ such that $(P^n)_{ij} > 0$.

Proof. Immediate from the definition: $i \to j$ iff $\exists n \geq 0$ with $p_{ij}^{(n)} > 0$, so irreducibility is equivalent to all entries of $\sum_{n=0}^{N} P^n$ being positive for some N, which for finite chains follows from the accessibility of all states. □

2.4 Periodicity

Definition: Period

The **period** of state i is

$$d(i) = \gcd\{n \geq 1 : p_{ii}^{(n)} > 0\}.$$

State i is **aperiodic** if $d(i) = 1$ and **periodic** if $d(i) > 1$. An irreducible chain is aperiodic if any (equivalently, all) of its states have period 1.

Theorem: Period is a Class Property

If $i \leftrightarrow j$, then $d(i) = d(j)$. Therefore, every state in a communication class has the same period.

Proof. Since $i \leftrightarrow j$, there exist $r, s \geq 1$ such that $p_{ij}^{(r)} > 0$ and $p_{ji}^{(s)} > 0$. Therefore $p_{ii}^{(r+s)} \geq p_{ij}^{(r)} p_{ji}^{(s)} > 0$, which implies $d(i) \mid (r+s)$.

For any n with $p_{jj}^{(n)} > 0$, we have $p_{ii}^{(r+n+s)} \geq p_{ij}^{(r)} p_{jj}^{(n)} p_{ji}^{(s)} > 0$, so $d(i) \mid (r+n+s)$. Combined with $d(i) \mid (r+s)$, we get $d(i) \mid n$ for all n in the return-time set of j. Hence $d(i) \mid d(j)$. Symmetrically $d(j) \mid d(i)$, so $d(i) = d(j)$. □

Example: Periodic Chain

Let $\mathcal{X} = \{0, 1, 2\}$ with transitions $0 \to 1 \to 2 \to 0$ (cyclic). The period of every state is $d = 3$, since $p_{00}^{(n)} > 0$ iff $3 \mid n$. The chain is irreducible but not aperiodic.

2.5 First-Passage Decompositions and the Green's Function

Before treating the strong Markov property in full generality, it is useful to develop the first-passage decomposition for return probabilities in an algebraically explicit form. This gives rise to the **Green's function** of the chain, a fundamental object connecting probability theory, potential theory, and the spectral theory developed in Part III.

Definition: Green's Function

The **Green's function** (or **potential kernel**) of a Markov chain with transition matrix P is the matrix

$$G(x,y) = \sum_{n=0}^{\infty} P^n(x,y) = \sum_{n=0}^{\infty} p_{xy}^{(n)},$$

counting the expected total number of visits to y starting from x. For a transient chain $G(x,y) < \infty$ for all x, y; for a recurrent chain $G(x,x) = \infty$ for every x.

The relationship between G and the first-passage probability $f_{xy}^{(n)} = \mathbb{P}_x(X_n = y, X_k \neq y \text{ for } 1 \le k < n)$ is given by the fundamental **renewal decomposition**:

Theorem: First-Passage Decomposition

For all $x, y \in \mathcal{X}$ and $n \ge 1$,

$$p_{xy}^{(n)} = \sum_{k=1}^{n} f_{xy}^{(k)}\, p_{yy}^{(n-k)}, \qquad p_{yy}^{(0)} = 1.$$

In generating-function form, writing $G_{xy}(z) = \sum_{n=0}^{\infty} p_{xy}^{(n)} z^n$ and $F_{xy}(z) = \sum_{n=1}^{\infty} f_{xy}^{(n)} z^n$:

$$G_{xy}(z) = \mathbf{1}_{x=y} + F_{xy}(z)\, G_{yy}(z).$$

Proof. The path from x to y in n steps must make its *first* visit to y at some time $k \in \{1, \dots, n\}$; by the Markov property applied at that first visit time,

$$p_{xy}^{(n)} = \sum_{k=1}^{n} \mathbb{P}_x(\text{first visit to } y \text{ at time } k) \cdot \mathbb{P}_y(X_{n-k} = y) = \sum_{k=1}^{n} f_{xy}^{(k)}\, p_{yy}^{(n-k)}.$$

Multiplying both sides by z^n and summing from $n = 1$ to ∞ gives the generating-function identity; the $x = y$ case contributes an extra 1 from the $n = 0$ term $p_{yy}^{(0)} = 1$. □

Corollary: Recurrence via Green's Function

State y is recurrent if and only if $G(y,y) = +\infty$. If y is transient, then $G(y,y) = (1 - F_{yy}(1))^{-1} < \infty$, where $F_{yy}(1) = \mathbb{P}_y(\tau_y < \infty) < 1$ is the return probability.

Proof. Setting $z = 1$ and $x = y$ in the generating-function identity: $G_{yy}(1) = 1 + F_{yy}(1)\, G_{yy}(1)$, which gives $G_{yy}(1)(1 - F_{yy}(1)) = 1$, hence $G_{yy}(1) = (1 - F_{yy}(1))^{-1}$. When $F_{yy}(1) = 1$ (recurrent), the formula gives ∞. □

Example: Green's Function for the Two-State Chain

For $P = \begin{pmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{pmatrix}$ on $\{1, 2\}$ with $\alpha, \beta > 0$, eigenvalues of P are 1 and

$1-\alpha-\beta$. The Green's function at $z<1$ is

$$G(z)=(I-zP)^{-1}=\frac{1}{(1-z)(1-z(1-\alpha-\beta))}\begin{pmatrix}1-z(1-\beta) & z\alpha\\ z\beta & 1-z(1-\alpha)\end{pmatrix}.$$

Taking $z\to 1^-$ confirms $G(1,1)=G(2,2)=\infty$, consistent with positive recurrence (as established in Chapter 3).

Proposition: Potential-Theoretic Interpretation

For a transient Markov chain, define the **potential operator** $G=\sum_{n=0}^{\infty}P^n$ (well-defined as a bounded operator on $\ell^\infty(\mathcal{X})$). A function $h:\mathcal{X}\to\mathbb{R}_+$ is **superharmonic** (i.e., $Ph\le h$ pointwise) if and only if h is an excessive function for the semigroup $(P^n)_{n\ge 0}$, and in this case G produces harmonic functions via $u=Gf$ for compactly supported f, generalizing the classical Newtonian potential.

This potential-theoretic perspective, while not pursued in full generality here, underpins the Foster–Lyapunov drift-condition approach to geometric ergodicity (Chapter 17) and the renewal-theoretic treatment of hitting times (Chapter 19): in both cases, carefully chosen Lyapunov/harmonic functions replace direct probabilistic estimates, reducing hard stochastic problems to the algebra of the operator G.

2.6 Stopping Times and the Strong Markov Property

The Markov property as stated is a fact about fixed, deterministic times n: conditional on X_n, the future is independent of the past. Many of the most useful arguments in the theory — first-passage decompositions, regeneration, the strong law for additive functionals — require the same independence to hold at a *random* time, provided that time is determined by the chain's own history rather than by foresight into the future. This is the content of the strong Markov property, and it rests on the notion of a stopping time.

Definition: Stopping Time

A random variable $T:\Omega\to\mathbb{N}_0\cup\{\infty\}$ is a **stopping time** (with respect to the natural filtration $\{\mathcal{F}_n\}$ of the chain) if

$$\{T\le n\}\in\mathcal{F}_n=\sigma(X_0,X_1,\dots,X_n)\qquad\text{for every } n\ge 0.$$

Equivalently, $\{T=n\}\in\mathcal{F}_n$ for every n: whether the chain has stopped by time n can be determined from the trajectory observed up to and including time n, with no knowledge of the future required.

Example: Hitting Times Are Stopping Times

The first hitting time of a set $A \subseteq \mathcal{X}$, defined by $T_A = \inf\{n \geq 0 : X_n \in A\}$ (with $T_A = \infty$ if the chain never visits A), is a stopping time, since $\{T_A \leq n\} = \bigcup_{k=0}^{n} \{X_k \in A\} \in \mathcal{F}_n$. By contrast, the *last* visit time to A before some horizon, or the time of the maximum of $\{X_0, \ldots, X_N\}$ read backward from the end, is in general *not* a stopping time, since determining it requires knowledge of the entire trajectory up to N, not merely up to the candidate time itself.

Theorem: Strong Markov Property

Let $\{X_n\}$ be a time-homogeneous Markov chain and let T be a stopping time with $\mathbb{P}(T < \infty) = 1$. Then, conditional on $T < \infty$ and $X_T = i$, the post-T process $\{X_{T+n}\}_{n \geq 0}$ is itself a Markov chain started at i with the same transition matrix P, and is independent of $\mathcal{F}_T = \sigma(X_0, \ldots, X_T)$.

Proof. Fix $i \in \mathcal{X}$ and a bounded function f on path space. It suffices to show

$$\mathbb{E}[f(X_{T+1}, X_{T+2}, \ldots) \,|\, \mathcal{F}_T, T < \infty, X_T = i] = \mathbb{E}_i[f(X_1, X_2, \ldots)],$$

where $\mathbb{E}_i$ denotes expectation for the chain started fresh at i. Decompose over the value of T:

$$\begin{aligned} \mathbb{E}[f(X_{T+1}, \ldots)\, \mathbf{1}_{T<\infty}\, \mathbf{1}_{X_T = i} \mid \mathcal{F}_T] &= \sum_{n=0}^{\infty} \mathbf{1}_{T=n}\, \mathbf{1}_{X_n = i}\, \mathbb{E}[f(X_{n+1}, \ldots) \mid \mathcal{F}_n] \\ &= \sum_{n=0}^{\infty} \mathbf{1}_{T=n}\, \mathbf{1}_{X_n = i}\, \mathbb{E}_i[f(X_1, X_2, \ldots)], \end{aligned}$$

where the second equality is the ordinary (weak) Markov property applied at the deterministic time n, using that $\{T = n\} \in \mathcal{F}_n$ so the indicator $\mathbf{1}_{T=n}$ is itself $\mathcal{F}_n$-measurable and may be pulled outside the conditional expectation along with $\mathbf{1}_{X_n=i}$. Since the right-hand side no longer depends on n, summing over n simply reconstructs $\mathbf{1}_{T<\infty}\mathbf{1}_{X_T=i}\, \mathbb{E}_i[f(X_1, \ldots)]$, which is the claimed identity. □

The strong Markov property is the rigorous justification for a technique used freely throughout this volume: the **first-step** or **first-passage decomposition**, in which one conditions on the value of the chain at a hitting time and restarts the analysis from there as if from a fresh initial condition. Without the strong Markov property, such restarts would be unjustified, since T depends on the chain's own randomness rather than being fixed in advance.

Example: Strong Markov Property in the Gambler's Ruin

Consider simple random walk on $\{0, 1, \ldots, N\}$ absorbed at 0 and N, and let $T = T_{\{0,N\}}$ be the absorption time, which is almost surely finite by recurrence of the walk on a finite interval. The strong Markov property guarantees that, having reached T, the (trivial, absorbed) future evolution depends only on $X_T \in \{0, N\}$ and not on the path taken to get there — this is exactly what licenses the classical

first-step recursion $h(i) = \frac{1}{2}h(i-1) + \frac{1}{2}h(i+1)$ for the ruin probability $h(i) = \mathbb{P}_i(X_T = N)$, since the recursion implicitly restarts the chain at $i \pm 1$ after one step, treating the future from that point as an independent copy of the original problem.

A particularly important corollary concerns the return times to a fixed state, which will be the basic objects of renewal theory in Chapter 19.

Corollary: Regeneration at Return Times

Let $\{X_n\}$ be irreducible and let $T_i^{(1)} < T_i^{(2)} < \cdots$ denote the successive return times to a fixed state i (so $T_i^{(1)} = \inf\{n \geq 1 : X_n = i\}$ and $T_i^{(k+1)} = \inf\{n > T_i^{(k)} : X_n = i\}$). Then the excursions

$$E_k = \left(X_{T_i^{(k)}}, X_{T_i^{(k)}+1}, \ldots, X_{T_i^{(k+1)}}\right), \qquad k \geq 1,$$

are i.i.d. as random elements of the path space, each distributed as the chain started at i and run until its first return to i.

Proof. Each $T_i^{(k)}$ is a stopping time (an inductive application of the hitting-time argument above, since $T_i^{(k+1)}$ is the first-hitting time of i for the shifted chain $\{X_{T_i^{(k)}+n}\}_{n\geq 0}$), and by the strong Markov property, conditional on $X_{T_i^{(k)}} = i$, the excursion E_k has the same law as a fresh excursion of the chain started at i, independent of $\mathcal{F}_{T_i^{(k)}}$ and hence of $E_1, \ldots, E_{k-1}$. Since this holds for every k, the excursions are mutually independent and identically distributed. □

This regeneration structure is the probabilistic heart of the renewal-reward approach to the ergodic theorem: because the excursions E_k are i.i.d., classical results for sums of independent random variables (the strong law of large numbers, the central limit theorem, Wald's identity) apply directly to functionals of the excursions, and translate — via the renewal-reward theorem developed in Chapter 19 — into long-run average and fluctuation results for the original chain.

Part II

Long-Run Behavior: Stationarity and Ergodicity

Chapter 3

Stationary Distributions

3.1 Definition and Existence

Definition: Stationary Distribution

A probability distribution π on $\mathcal{X}$ is **stationary** (or **invariant**) for the Markov chain with transition matrix P if

$$\pi P = \pi,$$

i.e. $\pi_j = \sum_{i \in \mathcal{X}} \pi_i \, p_{ij}$ for all $j \in \mathcal{X}$. Equivalently, π is a left eigenvector of P with eigenvalue 1.

If the chain starts in distribution π (i.e. $\mathbb{P}(X_0 = i) = \pi_i$), then at every subsequent time $\mathbb{P}(X_n = i) = \pi_i$: the distribution is preserved. The stationary distribution represents the long-run fraction of time the chain spends in each state.

Theorem: Existence for Finite Irreducible Chains

Every finite irreducible Markov chain has a unique stationary distribution π, with $\pi_i > 0$ for all $i \in \mathcal{X}$.

Proof. Existence. The matrix P has spectral radius 1 (since all row sums equal 1), and by the Perron–Frobenius theorem applied to a positive power of P (which is strictly positive by irreducibility), there exists a unique positive left eigenvector π up to scaling; normalizing gives $\sum_i \pi_i = 1$.

Uniqueness. Suppose μ is another stationary distribution. Write $\mu = \alpha\pi + (1-\alpha)\nu$ for some distribution ν and $\alpha \in [0,1]$. By stationarity of both μ and π, and the ergodic theorem (proved below), $\mu_i = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} p_{ji}^{(k)}$ for any j, which equals π_i. Hence $\mu = \pi$.

Positivity. $\pi_i > 0$ for all i follows from the strict positivity of the Perron eigenvector. $\square$

3.2 Detailed Balance and Reversibility

Definition: Reversibility

A Markov chain with transition matrix P and stationary distribution π is **reversible** (or **time-reversible**) if it satisfies **detailed balance**:

$$\pi_i \, p_{ij} = \pi_j \, p_{ji} \qquad \forall i, j \in \mathcal{X}.$$

Detailed balance implies stationarity: summing over i, $\sum_i \pi_i p_{ij} = \sum_i \pi_j p_{ji} = \pi_j \sum_i p_{ji} = \pi_j$, so $\pi P = \pi$. Reversibility is a stronger condition meaning the process looks statistically identical whether run forward or backward in time.

Theorem: Reversibility and Symmetry

A Markov chain is reversible with respect to π if and only if the operator P is self-adjoint in $L^2(\pi)$, i.e. $\langle Pf, g\rangle_\pi = \langle f, Pg\rangle_\pi$ for all $f, g \in L^2(\pi)$, where $\langle f, g\rangle_\pi = \sum_i f(i)g(i)\pi_i$.

Proof. $\langle Pf, g\rangle_\pi = \sum_{i,j} p_{ij} f(j) g(i) \pi_i$. Similarly $\langle f, Pg\rangle_\pi = \sum_{i,j} p_{ij} f(i) g(j) \pi_j$. These are equal for all f, g iff $\pi_i p_{ij} = \pi_j p_{ji}$ for all i, j, which is precisely detailed balance. □

3.3 Finding Stationary Distributions

For finite chains, π satisfies the linear system $\pi P = \pi$ with $\sum_i \pi_i = 1$. This is equivalent to $\pi(P - I) = 0$, a homogeneous system. Since $P - I$ has rank $|\mathcal{X}| - 1$ for irreducible chains, the solution space is one-dimensional; the unique probability vector in this space is π.

Example: Two-State Chain

Let $P = \begin{pmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{pmatrix}$ with $\alpha, \beta \in (0, 1)$. The stationary distribution satisfies $\pi_1 \alpha = \pi_2 \beta$ and $\pi_1 + \pi_2 = 1$, giving

$$\pi_1 = \frac{\beta}{\alpha + \beta}, \qquad \pi_2 = \frac{\alpha}{\alpha + \beta}.$$

When $\alpha = \beta$, we get $\pi = (1/2, 1/2)$, as expected by symmetry.

3.4 Transience and Recurrence

Definition: Recurrence and Transience

Define the **first return time** to state i as $\tau_i = \min\{n \geq 1 : X_n = i\}$ (with $\tau_i = \infty$ if the chain never returns). State i is:

- **recurrent** if $\mathbb{P}_i(\tau_i < \infty) = 1$;
- **transient** if $\mathbb{P}_i(\tau_i < \infty) < 1$;
- **positive recurrent** if $\mathbb{E}_i[\tau_i] < \infty$;
- **null recurrent** if $\mathbb{P}_i(\tau_i < \infty) = 1$ but $\mathbb{E}_i[\tau_i] = \infty$.

Theorem: Recurrence Criterion

State i is recurrent if and only if $\sum_{n=0}^{\infty} p_{ii}^{(n)} = \infty$.

Proof. Let $f_i = \mathbb{P}_i(\tau_i < \infty)$ be the return probability. By the renewal theory decomposition, $\sum_{n=0}^{\infty} p_{ii}^{(n)} = (1 - f_i)^{-1}$ when $f_i < 1$ (transient case, finite sum) and $= \infty$ when $f_i = 1$ (recurrent case). Formally: let $G_i = \mathbb{E}_i\left[\sum_{n=0}^{\infty} \mathbf{1}_{X_n = i}\right]$ count expected visits. On each visit, the chain restarts, so $G_i = 1 + f_i G_i$, giving $G_i = (1 - f_i)^{-1}$ for $f_i < 1$. □

Theorem: Positive Recurrence and Stationarity

For an irreducible Markov chain on a countable state space, the following are equivalent:

(i) The chain is positive recurrent;

(ii) The chain has a stationary distribution π;

(iii) $\pi_i = 1/\mathbb{E}_i[\tau_i] > 0$ for all i.

This theorem connects the dynamical property of positive recurrence to the existence of a normalizable stationary measure, unifying two different perspectives on long-run behavior.

3.5 Variance of Additive Functionals and the Asymptotic Variance Formula

The ergodic theorem guarantees almost-sure convergence of time averages but says nothing about the rate of fluctuations. For a quantitative understanding of estimation error in MCMC and renewal-type problems, one needs the asymptotic variance of $n^{-1/2} \sum_{k=0}^{n-1} f(X_k)$.

Definition: Asymptotic Variance

For an ergodic Markov chain with stationary distribution π and a function f with $\pi(f) = 0$, the **asymptotic variance** is

$$\sigma_f^2 = \lim_{n\to\infty} \frac{1}{n} \operatorname{Var}_\pi\left(\sum_{k=0}^{n-1} f(X_k)\right) = \operatorname{Var}_\pi(f) + 2\sum_{k=1}^{\infty} \operatorname{Cov}_\pi\big(f(X_0), f(X_k)\big).$$

Theorem: Spectral Formula for the Asymptotic Variance

For a reversible ergodic chain with eigenvalues $1 = \lambda_0 > \lambda_1 \geq \cdots \geq \lambda_{N-1}$ and $L^2(\pi)$-orthonormal eigenbasis $\{\phi_k\}$, writing $f = \sum_{k\geq 1} \hat{f}_k \phi_k$ (since $\pi(f) = 0$ implies $\hat{f}_0 = 0$):

$$\sigma_f^2 = \sum_{k\geq 1} \frac{1+\lambda_k}{1-\lambda_k}\, \hat{f}_k^2 = \langle f, (I+P)(I-P)^{-1} f\rangle_\pi.$$

Proof. Using stationarity and the spectral decomposition, $\operatorname{Cov}_\pi(f(X_0), f(X_k)) = \langle f, P^k f\rangle_\pi = \sum_{j\geq 1} \lambda_j^k \hat{f}_j^2$. Therefore

$$\sigma_f^2 = \sum_{j\geq 1} \hat{f}_j^2 + 2\sum_{k=1}^{\infty}\sum_{j\geq 1} \lambda_j^k \hat{f}_j^2 = \sum_{j\geq 1} \hat{f}_j^2\left(1 + \frac{2\lambda_j}{1-\lambda_j}\right) = \sum_{j\geq 1} \frac{1+\lambda_j}{1-\lambda_j}\hat{f}_j^2,$$

where the geometric series $\sum_{k=1}^{\infty} \lambda_j^k = \lambda_j/(1-\lambda_j)$ converges for $|\lambda_j| < 1$, which holds for all $j \geq 1$ in an irreducible aperiodic chain. □

Corollary: Asymptotic Variance Bound via Spectral Gap

For a reversible chain with spectral gap $\gamma = 1 - \lambda_1$:

$$\sigma_f^2 \leq \frac{2}{\gamma} \operatorname{Var}_\pi(f).$$

Chains with larger spectral gap produce more efficient estimators: the effective sample size in n steps is $n\gamma/2$ rather than n.

Example: Asymptotic Variance for the Two-State Chain

On $\{0,1\}$ with $P(0,1) = \alpha$, $P(1,0) = \beta$ and $\pi(1) = \alpha/(\alpha+\beta)$, take $f(x) = x - \pi(1)$. The second eigenvalue is $\lambda_1 = 1-\alpha-\beta$ and $\hat{f}_1^2 = \alpha\beta/(\alpha+\beta)^3$. Therefore

$$\sigma_f^2 = \frac{1+(1-\alpha-\beta)}{1-(1-\alpha-\beta)} \cdot \frac{\alpha\beta}{(\alpha+\beta)^3} = \frac{2-\alpha-\beta}{\alpha+\beta} \cdot \frac{\alpha\beta}{(\alpha+\beta)^2} = \frac{(2-\alpha-\beta)\alpha\beta}{(\alpha+\beta)^3}.$$

When $\alpha = \beta$ this simplifies to $\sigma_f^2 = (2-2\alpha)/(4\alpha)\cdot 1/4 = (1-\alpha)/(8\alpha)$, confirming that rapid mixing ($\alpha \to 1/2$) reduces the asymptotic variance as predicted by the spectral bound above.

Theorem: Poisson Equation and Variance Representation

The asymptotic variance admits the **Poisson equation representation**

$$\sigma_f^2 = 2\langle h, f\rangle_\pi - \mathrm{Var}_\pi(f) = \mathrm{Var}_\pi(h) - \mathrm{Var}_\pi(Ph),$$

where h is the unique (up to constants) solution of the **Poisson equation** $(I - P)h = f - \pi(f)$, also called the **fundamental solution** or the **centered potential**.

The Poisson equation plays a role in Markov chain theory analogous to that of the resolvent in functional analysis: the operator $(I - P)^{-1}$ (restricted to the orthogonal complement of constants in $L^2(\pi)$) inverts the generator and maps forcing functions f to harmonic corrections h. This representation is used in Chapter 23 (MDP bias functions), Chapter 19 (renewal-theoretic variance formulas), and Chapter 20 (large-deviation rate functions), showing the Poisson equation as a unifying device across the volume.

Chapter 4

Ergodic Theory for Markov Chains

4.1 The Ergodic Theorem

The ergodic theorem is the cornerstone of the long-run analysis: it tells us that time averages converge to space averages under the stationary distribution.

Theorem: Ergodic Theorem for Markov Chains

Let $\{X_n\}$ be an irreducible, positive recurrent Markov chain with stationary distribution π. Then for any bounded function $f : \mathcal{X} \to \mathbb{R}$ and any initial state $i \in \mathcal{X}$:

$$\frac{1}{n}\sum_{k=0}^{n-1} f(X_k) \xrightarrow{\mathbb{P}_i\text{-a.s.}} \sum_{j\in\mathcal{X}} f(j)\,\pi_j = \mathbb{E}_\pi[f] \quad \text{as } n \to \infty.$$

Proof sketch. By the renewal theorem, the empirical occupation measure $\frac{1}{n}\sum_{k=0}^{n-1}\mathbf{1}_{X_k=j} \to 1/\mathbb{E}_j[\tau_j] = \pi_j$ almost surely. Applying this with $f(j)$ and using boundedness to justify the interchange of limit and sum gives the result. □

4.2 Convergence to Stationarity

Beyond time averages, we want to understand how the distribution at time n approaches π. The standard metric for this is the **total variation distance**.

Definition: Total Variation Distance

The **total variation distance** between probability measures μ and ν on $(\mathcal{X}, 2^{\mathcal{X}})$ is

$$\|\mu - \nu\|_{\mathrm{TV}} = \max_{A\subseteq\mathcal{X}} |\mu(A) - \nu(A)| = \frac{1}{2}\sum_{i\in\mathcal{X}} |\mu_i - \nu_i|.$$

Theorem: Convergence Theorem

Let $\{X_n\}$ be an irreducible, aperiodic, positive recurrent Markov chain on a finite state space with stationary distribution π. Then for any initial distribution μ:

$$\lim_{n\to\infty} \|\mu P^n - \pi\|_{\mathrm{TV}} = 0.$$

The proof uses the Perron–Frobenius theorem and the coupling argument developed in Part IV. The rate of convergence is governed by the *spectral gap*, the subject of Part III.

Part III

Spectral Theory and the Spectral Gap

Chapter 5

Spectral Analysis of Transition Matrices

5.1 Eigenvalues of Stochastic Matrices

For a finite Markov chain with $|\mathcal{X}| = N$, the transition matrix $P \in \mathbb{R}^{N\times N}$ has eigenvalues $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N \geq -1$.

Theorem: Eigenvalue Bounds

All eigenvalues of a row-stochastic matrix P satisfy $|\lambda_i| \leq 1$. The eigenvalue $\lambda_1 = 1$ always exists (with left eigenvector $\mathbf{1}$). For an irreducible chain, $\lambda_1 = 1$ has multiplicity 1 (i.e. π is the unique stationary distribution).

Proof. Let λ be an eigenvalue with right eigenvector v, so $Pv = \lambda v$. Choose $i^* = \arg\max_i |v_i|$. Then $|\lambda||v_{i^*}| = |(\lambda v)_{i^*}| = |(Pv)_{i^*}| = \left|\sum_j p_{i^*j} v_j\right| \leq \sum_j p_{i^*j}|v_j| \leq |v_{i^*}| \sum_j p_{i^*j} = |v_{i^*}|$, so $|\lambda| \leq 1$. Multiplicity-1 of $\lambda = 1$ follows from the Perron–Frobenius theorem applied to the positive matrix P^n for sufficiently large n (using irreducibility). □

5.2 The Spectral Gap

Definition: Spectral Gap

For an irreducible, reversible Markov chain, the eigenvalues satisfy $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_N \geq -1$. The **absolute spectral gap** is

$$\gamma^* = 1 - \lambda_* \quad \text{where} \quad \lambda_* = \max(|\lambda_2|, |\lambda_N|).$$

The **spectral gap** (ignoring the -1 eigenvalue) is $\gamma = 1 - \lambda_2$.

For aperiodic chains, $\lambda_N > -1$, so $\lambda_* = \lambda_2$ and $\gamma^* = \gamma = 1 - \lambda_2$.

Theorem: Spectral Gap Bound on Mixing

For a reversible, irreducible, aperiodic chain with spectral gap γ and stationary distribution π:

$$\|P^n(x,\cdot) - \pi\|_{\mathrm{TV}} \le \sqrt{\frac{1}{\pi_{\min}}}\,(1-\gamma)^n,$$

where $\pi_{\min} = \min_i \pi_i$.

Proof. Using the spectral decomposition $P = \sum_{k=1}^{N} \lambda_k \phi_k \phi_k^T$ (in $L^2(\pi)$-orthonormal eigenvectors ϕ_k, with $\phi_1 = \mathbf{1}$):

$$P^n(x,y) - \pi_y = \pi_y \sum_{k=2}^{N} \lambda_k^n \phi_k(x)\phi_k(y).$$

By Cauchy–Schwarz and Parseval:

$$\begin{aligned}\|P^n(x,\cdot) - \pi\|_{\mathrm{TV}}^2 &\le \frac{1}{4}\sum_{y} \frac{(P^n(x,y)-\pi_y)^2}{\pi_y} \\ &= \frac{1}{4}\sum_{k\ge 2} \lambda_k^{2n}\phi_k(x)^2 \le \frac{\lambda_*^{2n}}{4}\sum_{k\ge 2}\phi_k(x)^2 \le \frac{(1-\gamma)^{2n}}{4\pi_x}.\end{aligned}$$

Taking the square root gives the result (with $\pi_{\min}$ bounding π_x from below uniformly). □

5.3 Canonical Paths and the Congestion Ratio

The **canonical paths** method, introduced by Sinclair and Jerrum, provides a systematic approach to proving lower bounds on conductance (and hence upper bounds on mixing time) by routing probability flow through the graph of the chain.

Definition: Canonical Paths

Let $\Gamma = \{\gamma_{xy}\}_{x,y\in\mathcal{X}}$ be a set of **canonical paths** where each γ_{xy} is a path in the directed graph of the chain from x to y. The **congestion ratio** is

$$\varrho(\Gamma) = \max_{(u,v):P(u,v)>0} \frac{1}{\pi(u)P(u,v)} \sum_{\substack{x,y\in\mathcal{X}:\\ (u,v)\in\gamma_{xy}}} \pi(x)\pi(y)|\gamma_{xy}|,$$

where $|\gamma_{xy}|$ denotes the length (number of edges) of γ_{xy}.

Theorem: Canonical Paths Bound (Sinclair–Jerrum)

For a reversible Markov chain,

$$\gamma \ge \frac{1}{\varrho(\Gamma)}.$$

Proof. The Dirichlet form can be written $\mathcal{E}(f,f) = \frac{1}{2}\sum_{u,v}\pi(u)P(u,v)(f(u)-f(v))^2$. For any f with $\pi(f)=0$, we have by writing $(f(x)-f(y))$ as a telescoping sum along γ_{xy}:

$$\begin{aligned}
\operatorname{Var}_\pi(f) &= \frac{1}{2}\sum_{x,y}\pi(x)\pi(y)(f(x)-f(y))^2 \\
&\le \frac{1}{2}\sum_{x,y}\pi(x)\pi(y)|\gamma_{xy}| \sum_{(u,v)\in\gamma_{xy}}(f(u)-f(v))^2 \\
&= \frac{1}{2}\sum_{(u,v)}(f(u)-f(v))^2 \sum_{\substack{x,y:\\ (u,v)\in\gamma_{xy}}}\pi(x)\pi(y)|\gamma_{xy}| \\
&\le \varrho(\Gamma)\cdot\mathcal{E}(f,f),
\end{aligned}$$

where the first inequality is Cauchy–Schwarz on the telescoping sum. The variational formula $\gamma = \inf_f \mathcal{E}(f,f)/\operatorname{Var}_\pi(f)$ then gives $\gamma \ge 1/\varrho(\Gamma)$. □

Example: Canonical Paths for the Hypercube

For the lazy random walk on $\{0,1\}^n$, route the canonical path from x to y by correcting each coordinate in turn (left to right): flip bit 1 first, then bit 2, etc. Each path has length at most n. An edge (u,v) (differing in coordinate i) is used by a path γ_{xy} whenever x and y agree on bits $1,\dots,i-1$ and differ on bit i, giving at most 2^n such pairs per edge. Therefore

$$\varrho(\Gamma) \le n\cdot\frac{2^n\cdot(1/2^n)(1/2^n)}{(1/2^n)(1/(2n))} = 2n^2,$$

so $\gamma \ge 1/(2n^2)$ and $t_{\text{mix}} = O(n^2\log n)$. The spectral gap approach (Chapter 5) gives the sharper bound $\gamma = 1/n$ and $t_{\text{mix}} = O(n\log n)$; the canonical paths method gives a weaker but more combinatorially transparent bound, and its power is greatest for more complex chains (Markov chains on combinatorial structures like graphs, matchings, or colorings) where direct spectral analysis is intractable.

Remark: Comparison Theorem

The canonical paths technique can also be used to *compare* two different chains on the same state space. If chain $\tilde{P}$ has known spectral gap $\tilde{\gamma}$ and one can route canonical paths of P through edges of $\tilde{P}$ with congestion A, then $\gamma(P) \ge \tilde{\gamma}/A$. This **comparison theorem** (Diaconis–Saloff-Coste) is particularly useful for analyzing Gibbs samplers and other structured MCMC chains by comparison to simpler reference chains.

Definition: Dirichlet Form

For a reversible chain with transition matrix P and stationary distribution π, the **Dirichlet form** is

$$\mathcal{E}(f,f) = \frac{1}{2}\sum_{i,j} \pi_i p_{ij}(f(j)-f(i))^2 = \langle f, (I-P)f\rangle_\pi.$$

Theorem: Variational Formula for Spectral Gap

The spectral gap satisfies

$$\gamma = \inf_{f \neq \text{const}} \frac{\mathcal{E}(f,f)}{\mathrm{Var}_\pi(f)},$$

where $\mathrm{Var}_\pi(f) = \sum_i \pi_i f(i)^2 - \left(\sum_i \pi_i f(i)\right)^2$.

This variational characterization — the **Poincaré inequality** — is the foundation for bounding the spectral gap from below using geometric properties of the state space and the graph structure of the chain.

5.4 Cheeger's Inequality

The **Cheeger constant** (or **conductance**) of a Markov chain connects the geometric structure of the state space to the spectral gap.

Definition: Conductance

For a Markov chain with transition matrix P and stationary distribution π, the **conductance** of a set $S \subseteq \mathcal{X}$ with $0 < \pi(S) \leq 1/2$ is

$$\Phi(S) = \frac{\sum_{i\in S, j\notin S} \pi_i p_{ij}}{\pi(S)}.$$

The **Cheeger constant** is $\Phi_* = \min_{S: 0<\pi(S)\leq 1/2} \Phi(S)$.

Theorem: Cheeger's Inequality

For a reversible Markov chain:

$$\frac{\Phi_*^2}{2} \leq \gamma \leq 2\Phi_*.$$

The lower bound shows that poor conductance (bottlenecks in the chain) implies a small spectral gap and hence slow mixing. The upper bound shows that conductance completely controls the gap up to constant factors.

Part IV

Coupling Methods and Mixing Time Bounds

Chapter 6

The Coupling Technique

6.1 Definition and Basic Properties

Definition: Coupling

A **coupling** of probability measures μ and ν on $\mathcal{X}$ is a pair of random variables (X, Y) on a common probability space such that $X \sim \mu$ and $Y \sim \nu$ (i.e. their marginals are μ and ν respectively).

The coupling inequality is the fundamental tool connecting couplings to total variation distance.

Theorem: Coupling Inequality

For any coupling (X, Y) of μ and ν:

$$\|\mu - \nu\|_{\mathrm{TV}} \leq \mathbb{P}(X \neq Y).$$

Equality holds for the **optimal coupling**, which achieves $\mathbb{P}(X \neq Y) = \|\mu - \nu\|_{\mathrm{TV}}$.

Proof. For any $A \subseteq \mathcal{X}$:

$$\mu(A) - \nu(A) = \mathbb{P}(X \in A) - \mathbb{P}(Y \in A) \leq \mathbb{P}(X \in A, Y \notin A) \leq \mathbb{P}(X \neq Y).$$

Taking the maximum over A gives $\|\mu - \nu\|_{\mathrm{TV}} \leq \mathbb{P}(X \neq Y)$. For the optimal coupling, set $A^* = \{i : \mu_i > \nu_i\}$ and construct (X, Y) to agree on A^* and $(A^*)^c$ with probability $\nu(A^*) = \mu(A^{*c})$, so $\mathbb{P}(X \neq Y) = \mu(A^*) - \nu(A^*) = \|\mu - \nu\|_{\mathrm{TV}}$. □

6.2 Markovian Couplings

A **Markovian coupling** of a Markov chain with transition matrix P started from states x and y is a joint process $\{(X_n, Y_n)\}_{n \geq 0}$ on $\mathcal{X} \times \mathcal{X}$ where both $\{X_n\}$ and $\{Y_n\}$ are

individually Markov chains with transition matrix P. The **coupling time** is $\tau_{\text{couple}} = \min\{n \geq 0 : X_n = Y_n\}$.

After coupling (once $X_n = Y_n$), we can run the two chains together ($X_{n+k} = Y_{n+k}$ for all $k \geq 0$), using the same randomness. This is called **coalescence**.

Theorem: Coupling Bound on Mixing

For any Markovian coupling of a chain starting at x and from stationarity π:

$$\|P^n(x,\cdot) - \pi\|_{\text{TV}} \leq \mathbb{P}_x(\tau_{\text{couple}} > n).$$

6.3 The Mixing Time

Definition: Mixing Time

The ε**-mixing time** of a Markov chain is

$$t_{\text{mix}}(\varepsilon) = \min\left\{n \geq 0 : \max_{x \in \mathcal{X}} \|P^n(x,\cdot) - \pi\|_{\text{TV}} \leq \varepsilon\right\}.$$

The standard mixing time is $t_{\text{mix}} = t_{\text{mix}}(1/4)$.

Theorem: Mixing Time and Spectral Gap

For a reversible, irreducible, aperiodic chain:

$$\left(\frac{1}{\gamma} - 1\right) \ln\frac{1}{2\varepsilon} \leq t_{\text{mix}}(\varepsilon) \leq \frac{1}{\gamma} \ln\frac{1}{\varepsilon\, \pi_{\min}}.$$

The lower and upper bounds show that $t_{\text{mix}} \asymp \gamma^{-1}$ up to logarithmic factors in $\pi_{\min}$, confirming the spectral gap as the key quantity governing convergence speed.

6.4 Doeblin's Coupling and Quantitative Convergence

A particularly elegant and historically important coupling is **Doeblin's coupling**, which applies whenever the chain satisfies a *one-step minorization* over the entire state space — the classical **Doeblin condition**.

Definition: Doeblin Condition

A Markov chain with transition matrix P satisfies the **Doeblin condition** if there exist $\varepsilon > 0$, $n_0 \geq 1$, and a probability measure ν on $\mathcal{X}$ such that

$$P^{n_0}(x,\cdot) \geq \varepsilon\, \nu(\cdot) \qquad \text{for all } x \in \mathcal{X}.$$

Theorem: Doeblin's Coupling Theorem

If the Doeblin condition holds with constants (ε, n_0, ν), then for any two initial distributions μ_1, μ_2 and any $n \geq 0$:

$$\|\mu_1 P^{n_0 n} - \mu_2 P^{n_0 n}\|_{\mathrm{TV}} \leq (1-\varepsilon)^n.$$

In particular, $t_{\mathrm{mix}}(\varepsilon') \leq n_0 \lceil \log(\varepsilon')^{-1} / \log(1-\varepsilon)^{-1} \rceil$ for any $\varepsilon' > 0$.

Proof. Construct a coupling of the two chains as follows. At each block of n_0 steps, with probability ε draw the next state from ν (the same draw for both chains); with probability $1-\varepsilon$, evolve each chain independently according to its respective distribution. The coupling succeeds (both chains hit the same state) with probability ε in each block, independently across blocks. After n blocks, the probability of not yet coupling is $(1-\varepsilon)^n$. Applying the coupling inequality gives the stated total-variation bound. □

Proposition: Doeblin Condition Implies Exponential Ergodicity

Under the Doeblin condition, the chain is uniformly geometrically ergodic: there exist $C < \infty$ and $r \in (0,1)$ such that

$$\sup_{x \in \mathcal{X}} \|P^n(x, \cdot) - \pi\|_{\mathrm{TV}} \leq C r^n \qquad \text{for all } n \geq 0.$$

The Doeblin condition is a global version of the minorization condition used in Chapter 17 for general state spaces. For finite state spaces, every irreducible aperiodic chain satisfies Doeblin's condition (since P^{n_0} has all positive entries for sufficiently large n_0, by irreducibility and aperiodicity), which provides an alternative route to the basic convergence theorem of Part II.

Example: Doeblin Coupling for Random Walk on K_n

For simple random walk on the complete graph K_n (transition: move to a uniformly random other vertex), $P(x,y) = 1/(n-1)$ for $x \neq y$. Taking $\nu = \pi$ (the uniform distribution), $P(x,\cdot) \geq \frac{1}{n-1}\pi(\cdot) \cdot (n-1) \cdot \mathbf{1}_{\{y \neq x\}}$ is not quite of Doeblin form directly, but $P^2(x,\cdot) \geq \frac{1}{n-1}\pi(\cdot)$ for all x, giving $\varepsilon = 1/(n-1)$ and $n_0 = 2$. Doeblin's theorem then yields $t_{\mathrm{mix}} \leq O(n \log n)$, matching the exact mixing time $t_{\mathrm{mix}} = O(\log n)$ only up to a polynomial factor; sharper spectral arguments from Chapter 5 (eigenvalues: $\lambda_1 = -1/(n-1)$, $\gamma = n/(n-1)$) give the tight bound $t_{\mathrm{mix}}(1/4) = O(\log n)$.

Remark: Historical Note

Wolfgang Doeblin proved the coupling theorem bearing his name in the late 1930s, predating by decades the systematic development of coupling methods by Lindvall, Pitman, and Thorisson. His handwritten manuscript, sealed in an envelope and deposited with the French Academy of Sciences in 1940 (shortly before his death in World War II), was opened only in 2000 and revealed deep

insights into Brownian motion that anticipated much of later stochastic calculus. Doeblin's coupling remains, to this day, one of the most elegant and broadly applicable tools in the theory.

Path coupling is a powerful technique for bounding coupling times without constructing an explicit global coupling. The idea is to define a coupling for adjacent states in a graph and then extend to all pairs via paths.

Theorem: Path Coupling (Bubley–Dyer)

Let $G = (\mathcal{X}, \mathcal{E})$ be a connected graph with integer-valued path metric d. Suppose there exists a Markovian coupling such that for every edge $(u, v) \in \mathcal{E}$ (with $d(u, v) = 1$):

$$\mathbb{E}[d(X_1, Y_1) \mid X_0 = u, Y_0 = v] \leq \alpha$$

for some $\alpha < 1$. Then $t_{\text{mix}}(\varepsilon) \leq \frac{\ln(\operatorname{diam}(G)/\varepsilon)}{\ln(1/\alpha)}$.

Example: Random Transposition Walk

On $\mathcal{X} = S_n$ (permutations of $\{1, \ldots, n\}$), the random transposition walk picks a random transposition (i, j) uniformly and applies it. The mixing time is $t_{\text{mix}} \asymp n \ln n$. This was proved by Diaconis and Shahshahani using representation theory, and the path coupling approach provides an alternative route via the Hamming metric on permutations.

Chapter 7

Advanced Coupling Techniques

7.1 Monotone Coupling and Stochastic Ordering

When the state space $\mathcal{X}$ carries a partial order $\preceq$, one can often construct **monotone couplings** that preserve the ordering. This is particularly useful for Markov chains arising in statistical physics (e.g. the Ising model).

Definition: Stochastic Dominance

Measure μ **stochastically dominates** ν (written $\mu \succeq \nu$) if $\mu(f) \geq \nu(f)$ for every increasing function $f : \mathcal{X} \to \mathbb{R}$. For real-valued random variables, $X \succeq Y$ iff $\mathbb{P}(X > t) \geq \mathbb{P}(Y > t)$ for all t.

Theorem: Strassen's Theorem

$\mu \succeq \nu$ if and only if there exists a coupling (X, Y) of μ and ν with $X \geq Y$ almost surely.

7.2 Coupling from the Past

Coupling from the past (CFTP), introduced by Propp and Wilson, is a remarkable algorithm that produces exact samples from the stationary distribution π rather than merely approximate ones.

Algorithm: Propp–Wilson CFTP

1. Fix a time horizon T (starting with $T = 1$ and doubling).
2. Generate a single random mapping $\phi_t : \mathcal{X} \to \mathcal{X}$ for each $t \in \{-T+1, \ldots, 0\}$ using common random variables across all starting states.
3. Compute $X_0^{(x)} = \phi_0 \circ \phi_{-1} \circ \cdots \circ \phi_{-T+1}(x)$ for every $x \in \mathcal{X}$.
4. If $X_0^{(x)}$ is the same for all x (coalescence), output the common value. Oth-

erwise, double T and restart.

The output has exactly the distribution π.

The correctness of CFTP relies on the fact that, once all chains started from different initial states coalesce to a single value at time 0, that value is exactly distributed as π, independent of the starting state.

Part V

Continuous-Time Markov Chains

Chapter 8

Generators and Semigroups

8.1 Definition of Continuous-Time Markov Chains

Definition: Continuous-Time Markov Chain

A stochastic process $\{X_t\}_{t\geq 0}$ on a countable state space $\mathcal{X}$ is a **continuous-time Markov chain** (CTMC) if it satisfies the Markov property: for all $s, t \geq 0$ and measurable $A \subseteq \mathcal{X}$,

$$\mathbb{P}(X_{t+s} \in A \mid \mathcal{F}_t) = \mathbb{P}(X_{t+s} \in A \mid X_t).$$

A CTMC is **time-homogeneous** if $\mathbb{P}(X_{t+s} \in A \mid X_t = i) = \mathbb{P}(X_s \in A \mid X_0 = i)$.

The key structural result is that a continuous-time Markov chain with right-continuous paths can be decomposed as follows: the chain remains in state i for a time $\mathrm{Exp}(q_i)$-distributed holding time (with rate $q_i = \sum_{j\neq i} q_{ij}$), then jumps to state $j \neq i$ with probability q_{ij}/q_i.

8.2 The Generator

Definition: Generator

The **generator** (or **Q-matrix**, or **rate matrix**) of a time-homogeneous CTMC is the matrix $Q = (q_{ij})_{i,j\in\mathcal{X}}$ with:

$$q_{ij} = \lim_{t\downarrow 0} \frac{p_{ij}(t) - \delta_{ij}}{t} \geq 0 \quad (i \neq j), \qquad q_{ii} = -\sum_{j\neq i} q_{ij} \leq 0.$$

Each row of Q sums to 0: $\sum_j q_{ij} = 0$.

The transition semigroup $P(t) = (p_{ij}(t))$ satisfies the **Kolmogorov forward and backward equations**:

$$\frac{d}{dt}P(t) = P(t)Q \quad \text{(forward)}, \qquad \frac{d}{dt}P(t) = QP(t) \quad \text{(backward)}.$$

The solution is the **matrix exponential**: $P(t) = e^{tQ} = \sum_{n=0}^{\infty} \frac{(tQ)^n}{n!}$.

8.3 Stationary Distribution and Reversibility

A distribution π is stationary for a CTMC iff $\pi Q = 0$. The detailed balance condition for CTMCs is $\pi_i q_{ij} = \pi_j q_{ji}$ for all $i \neq j$. Reversibility and spectral analysis carry over from the discrete-time theory, with the generator Q replacing $I - P$.

Example: Birth-Death Process

A **birth-death process** on $\mathcal{X} = \{0, 1, 2, \ldots\}$ has $q_{i,i+1} = \lambda_i$ (birth rates) and $q_{i,i-1} = \mu_i$ (death rates) with all other off-diagonal rates zero. It is always reversible. The stationary distribution (when it exists) satisfies the detailed balance recursion $\pi_{i+1} = \frac{\lambda_i}{\mu_{i+1}} \pi_i$, giving

$$\pi_n = \pi_0 \prod_{k=0}^{n-1} \frac{\lambda_k}{\mu_{k+1}}.$$

Birth-death processes model queues (Erlang's M/M/1), population dynamics, and chemical reactions.

Chapter 9

Reversible Chains and Spectral Methods in Continuous Time

9.1 The Relaxation Time

For a reversible CTMC with generator Q, the eigenvalues of $-Q$ are $0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \cdots$. The **spectral gap** is λ_1 and the **relaxation time** is

$$t_{\mathrm{rel}} = \frac{1}{\lambda_1}.$$

Theorem: Exponential Convergence

For a reversible, irreducible CTMC:

$$\|P_x(t) - \pi\|_{\mathrm{TV}} \le \frac{1}{2}\sqrt{\frac{1}{\pi_{\min}}}\, e^{-\lambda_1 t}.$$

The mixing time satisfies $t_{\mathrm{mix}}(\varepsilon) \le t_{\mathrm{rel}} \ln\!\left(\frac{1}{2\varepsilon\sqrt{\pi_{\min}}}\right)$.

9.2 Log-Sobolev Inequalities

The **logarithmic Sobolev inequality** provides sharper mixing bounds than Poincaré alone.

Definition: Log-Sobolev Constant

The **log-Sobolev constant** ρ is the largest constant such that for all $f > 0$:

$$\mathrm{Ent}_\pi(f^2) \le \frac{2}{\rho}\,\mathcal{E}(f, f),$$

where $\mathrm{Ent}_\pi(f) = \mathbb{E}_\pi[f \ln f] - \mathbb{E}_\pi[f] \ln \mathbb{E}_\pi[f]$ is the entropy of f under π.

Theorem: Log-Sobolev Mixing Bound

If the log-Sobolev inequality holds with constant ρ, then

$$t_{\text{mix}}(\varepsilon) \leq \frac{1}{2\rho}\left(\ln\ln\frac{1}{\pi_{\min}} + \ln\frac{1}{2\varepsilon^2}\right).$$

This improves over the Poincaré bound by removing the dependence on $\ln(1/\pi_{\min})$ in many cases.

Part VI

Applications: MCMC, Random Walks, and Beyond

Chapter 10

Markov Chain Monte Carlo

10.1 The Metropolis–Hastings Algorithm

Markov Chain Monte Carlo (MCMC) methods construct Markov chains whose stationary distribution is a target π that is known up to a normalizing constant. The Metropolis–Hastings algorithm is the foundational MCMC method.

Algorithm: Metropolis–Hastings

Given a target π and a symmetric proposal kernel $Q(x, \cdot)$:

1. At state $X_n = x$, propose $Y \sim Q(x, \cdot)$.
2. Compute acceptance ratio $\alpha(x, y) = \min\left(1, \dfrac{\pi(y)Q(y, x)}{\pi(x)Q(x, y)}\right)$.
3. Set $X_{n+1} = Y$ with probability $\alpha(x, y)$; else $X_{n+1} = x$.

The chain has stationary distribution π by detailed balance.

Theorem: Detailed Balance for Metropolis–Hastings

The Metropolis–Hastings chain satisfies $\pi(x)P(x, y) = \pi(y)P(y, x)$ for all $x \neq y$.

Proof. For $x \neq y$, the transition density is $P(x, y) = Q(x, y)\alpha(x, y)$. Then:

$$\begin{aligned}
\pi(x)P(x, y) &= \pi(x)Q(x, y)\min\left(1, \frac{\pi(y)Q(y, x)}{\pi(x)Q(x, y)}\right) \\
&= \min(\pi(x)Q(x, y),\ \pi(y)Q(y, x)) \\
&= \pi(y)Q(y, x)\min\left(\frac{\pi(x)Q(x, y)}{\pi(y)Q(y, x)}, 1\right) = \pi(y)P(y, x).
\end{aligned}$$

□

10.2 Gibbs Sampling

Gibbs sampling is an MCMC method tailored to distributions over product spaces $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_d$.

Algorithm: Gibbs Sampler

At each step, choose a coordinate $k \in \{1, \dots, d\}$ (uniformly at random or in systematic scan) and resample:

$$X_k \sim \pi(\cdot \mid X_1, \dots, X_{k-1}, X_{k+1}, \dots, X_d).$$

The full conditionals $\pi(x_k \mid x_{-k})$ are often easy to compute even when π itself is not tractable.

Gibbs sampling is a special case of Metropolis–Hastings with acceptance probability $\alpha \equiv 1$. It is widely used in Bayesian statistics, latent variable models, and statistical physics.

10.3 Mixing Times for MCMC

The practical performance of MCMC depends critically on mixing times. For the random-walk Metropolis on $\mathbb{R}^d$ with a Gaussian target $\pi = \mathcal{N}(0, I_d)$ and step size σ:

- Optimal scaling: $\sigma^2 \approx 2.38^2/d$ (Roberts–Gelman–Gilks, 1997)
- Mixing time: $t_{\text{mix}} = O(d)$ at optimal scaling
- Acceptance rate at optimal scaling: $\approx 23.4\%$

These results, proved using diffusion limits (the chain converges weakly to a Langevin diffusion in high dimensions), are cornerstones of practical MCMC calibration.

Chapter 11

Random Walks on Graphs

11.1 The Lazy Random Walk

Let $G = (V, E)$ be a connected undirected graph with $|V| = N$. The **simple random walk** on G has transition matrix $P_{uv} = 1/\deg(u)$ for each edge $\{u, v\} \in E$. The **lazy random walk** sets $P_{uu} = 1/2$ and $P_{uv} = 1/(2\deg(u))$, making the chain aperiodic.

Theorem: Stationary Distribution of Random Walk

The stationary distribution of the simple random walk on G is $\pi_v = \deg(v)/(2|E|)$.

11.2 Spectral Gap for Random Walks

The spectral gap of the random walk on G is related to the combinatorial structure of G through the **graph Laplacian** $L = D - A$, where $D = \operatorname{diag}(\deg(v))$ and A is the adjacency matrix.

The normalized Laplacian is $\mathcal{L} = D^{-1/2} L D^{-1/2}$, with eigenvalues $0 = \mu_0 \le \mu_1 \le \cdots \le \mu_{N-1} \le 2$. The spectral gap of the lazy random walk is $\gamma = \mu_1/2$.

Example: Hypercube Random Walk

The d-dimensional hypercube $Q_d = \{0,1\}^d$ is a d-regular graph on 2^d vertices. The eigenvalues of the adjacency matrix are $\lambda_k = d - 2k$ for $k = 0, 1, \ldots, d$ with multiplicity $\binom{d}{k}$. The spectral gap of the lazy walk is $\gamma = 1/d$, giving $t_{\text{mix}} = O(d \log d)$. This is tight: the mixing time is $\Theta(d \log d)$.

11.3 Expander Graphs

Definition: Expander Graph

An infinite family of d-regular graphs $\{G_n\}$ with $|V(G_n)| \to \infty$ is an **expander family** if there exists $\gamma > 0$ such that the spectral gap of the random walk on G_n is at least γ for all n.

Expanders have $t_{\text{mix}} = O(\log N)$, the best possible for regular graphs. The existence of explicit expander families was proved by Margulis (1973) using Kazhdan's property T, and by Lubotzky–Phillips–Sarnak (1988) via Ramanujan graphs, where $\lambda_* \le 2\sqrt{d-1}$ (the Alon–Boppana bound).

Chapter 12

Birth-Death Chains and Queueing Theory

12.1 Spectral Theory of Birth-Death Chains

Birth-death chains have a particularly tractable spectral theory due to their tridiagonal structure.

Theorem: Spectral Gap of Birth-Death Chains

For an irreducible birth-death chain on $\{0, 1, \ldots, N\}$ with stationary distribution π, the spectral gap satisfies

$$\gamma \geq \frac{1}{\sum_{k=0}^{N} \frac{1}{\pi_k q_k}} \cdot \frac{1}{\max_k \pi_k}$$

where $q_k = p_k + r_k$ are the rates from state k.

12.2 The M/M/1 Queue

The M/M/1 queue is a continuous-time birth-death process on $\mathcal{X} = \{0, 1, 2, \ldots\}$ with arrival rate λ and service rate μ.

Theorem: M/M/1 Stationary Distribution

When $\rho = \lambda/\mu < 1$ (stability condition), the M/M/1 queue has unique stationary distribution $\pi_n = (1-\rho)\rho^n$, $n \geq 0$. The mean queue length is $\mathbb{E}_\pi[X] = \rho/(1-\rho)$ and the mean waiting time (Little's law) is $W = 1/(\mu - \lambda)$.

Proof. Detailed balance gives $\pi_n \lambda = \pi_{n+1}\mu$, so $\pi_{n+1} = \rho\,\pi_n$. This geometric recursion gives $\pi_n = \pi_0 \rho^n$. Normalization: $\pi_0 \sum_{n=0}^{\infty} \rho^n = \pi_0/(1-\rho) = 1$, so $\pi_0 = 1 - \rho$. $\square$

12.3 Hitting Times and Fundamental Matrix

For a Markov chain on finite state space $\mathcal{X}$, the **hitting time** to set $A \subseteq \mathcal{X}$ is $\tau_A = \min\{n \geq 0 : X_n \in A\}$. The expected hitting times $h_i = \mathbb{E}_i[\tau_A]$ satisfy the **Dirichlet problem**:

$$h_i = 0 \quad (i \in A), \qquad h_i = 1 + \sum_{j \notin A} p_{ij} h_j \quad (i \notin A).$$

This linear system has a unique solution, expressible via the **fundamental matrix** $Z = (I - (I - \Pi)P)^{-1}$ where $\Pi_{ij} = \pi_j$, giving $h_i = \sum_{j \notin A} Z_{ij} \cdot (\text{appropriate corrections})$.

Chapter 13

Convergence Rates: Lower Bounds and Cutoff Phenomenon

13.1 Lower Bounds on Mixing Times

Upper bounds on mixing times are obtained from couplings and spectral gaps. Lower bounds require different arguments. The main techniques are: the eigenvalue lower bound, the bottleneck lower bound via conductance, and the distinguishing statistic method.

Theorem: Eigenvalue Lower Bound

For a reversible chain with stationary distribution π and second-largest eigenvalue λ_2:

$$t_{\text{mix}}(\varepsilon) \geq \frac{1-2\varepsilon}{1-\lambda_2}.$$

In particular, $t_{\text{mix}} \geq t_{\text{rel}}/2 = 1/(2\gamma)$ up to constant factors.

Proof. Let $f = \phi_2$ be the second eigenfunction normalized so $\|f\|_{L^2(\pi)} = 1$ and $\mathbb{E}_\pi[f] = 0$. Starting from the state x_0 maximizing $|\phi_2(x_0)|/\sqrt{\pi_{x_0}}$:

$$\|P^n(x_0,\cdot) - \pi\|_{\text{TV}} \geq \frac{1}{2}|\mathbb{E}_{P^n(x_0,\cdot)}[f] - \mathbb{E}_\pi[f]| = \frac{1}{2}|P^n f(x_0)| = \frac{1}{2}\lambda_2^n |f(x_0)|.$$

Setting this $\geq \varepsilon$ gives $\lambda_2^n \geq 2\varepsilon/|f(x_0)| \geq 2\varepsilon$. Since $\lambda_2^n = (1-\gamma)^n \geq 1 - n\gamma$, we need $1 - n\gamma \geq 2\varepsilon$, giving $n \leq (1-2\varepsilon)/\gamma$. □

Theorem: Bottleneck Lower Bound

For any set $S \subseteq \mathcal{X}$ with $\pi(S) \leq 1/2$:

$$t_{\text{mix}}(\varepsilon) \geq \frac{1-2\varepsilon}{4\Phi(S)}.$$

Proof. Starting from the worst-case $x \in S$, the probability of being outside S increases by at most $\Phi(S)$ per step (the conductance bounds the probability flux). Formally,

$P^n(x, S^c) \le n\Phi(S)$, while $\pi(S^c) \ge 1/2$. So $\|P^n(x,\cdot) - \pi\|_{\mathrm{TV}} \ge \pi(S^c) - P^n(x,S^c) \ge 1/2 - n\Phi(S)$. Setting this $\ge \varepsilon$ gives the result. $\square$

13.2 The Cutoff Phenomenon

The **cutoff phenomenon** is one of the most striking results in the theory of mixing times: for many natural families of Markov chains, the transition from "far from stationarity" to "close to stationarity" occurs abruptly within a window much smaller than the mixing time itself.

Definition: Cutoff

A family of Markov chains $\{(P_n, \mathcal{X}_n, \pi_n)\}$ exhibits a **cutoff** if:

$$\lim_{n\to\infty} \frac{t_{\mathrm{mix}}^{(n)}(\varepsilon)}{t_{\mathrm{mix}}^{(n)}(1-\varepsilon)} = 1 \qquad \forall \varepsilon \in (0,1).$$

The **cutoff time** is t_n^* and the **window** is $w_n = o(t_n^*)$.

Example: Cutoff for Card Shuffling

The **random transposition shuffle** on n cards exhibits a cutoff at time $t_n^* = \frac{1}{2} n \ln n$ with window $w_n = O(n)$. Diaconis and Shahshahani showed:

$$\left\| P^t - \pi_{\mathrm{uniform}} \right\|_{\mathrm{TV}} \to \begin{cases} 1 & \text{if } t < \frac{1}{2}(1-\varepsilon) n \ln n \\ 0 & \text{if } t > \frac{1}{2}(1+\varepsilon) n \ln n \end{cases}$$

as $n \to \infty$ for any fixed $\varepsilon > 0$. The proof uses the representation theory of S_n.

Example: Cutoff for the Hypercube Walk

The lazy random walk on the n-dimensional hypercube $\{0,1\}^n$ exhibits a cutoff at time $t_n^* = \frac{1}{4} n \ln n$ with window $w_n = O(n)$. The total variation profile converges to a universal Gaussian error function profile as $n \to \infty$.

13.3 Product Chains and Tensorization

If $P_1, \dots, P_d$ are Markov chains on spaces $\mathcal{X}_1, \dots, \mathcal{X}_d$, their **product chain** evolves by choosing coordinate k at random and applying P_k. The spectral gap satisfies $\gamma_{\mathrm{prod}} = \min_k \gamma_k / d$.

Theorem: Tensorization of Entropy

If the log-Sobolev inequality with constant ρ_k holds for each factor P_k, then the product chain satisfies the log-Sobolev inequality with constant $\rho = \min_k \rho_k$.

The mixing time is $t_{\text{mix}}^{\text{prod}} = O\Big(d \cdot \max_k t_{\text{mix}}^{(k)}\Big)$.

Chapter 14

Martingales and Optional Stopping

14.1 Martingales

Definition: Martingale

A process $\{M_n\}_{n\geq 0}$ adapted to $\{\mathcal{F}_n\}$ is a **martingale** if $\mathbb{E}[|M_n|] < \infty$ and:

$$\mathbb{E}[M_{n+1} \mid \mathcal{F}_n] = M_n \quad \mathbb{P}\text{-a.s.}$$

Example: Harmonic Functions

Let $h : \mathcal{X} \to \mathbb{R}$ be harmonic for the chain: $Ph = h$. Then $\{h(X_n)\}$ is a martingale. The Dirichlet problem asks for harmonic h with boundary values $g|_A$; the solution is $h(i) = \mathbb{E}_i[g(X_{\tau_A})]$.

Theorem: Optional Stopping

Let $\{M_n\}$ be a martingale and τ a stopping time with $\mathbb{E}[\tau] < \infty$ and $|M_{n+1} - M_n| \leq C$ a.s. Then $\mathbb{E}[M_\tau] = \mathbb{E}[M_0]$.

Application: Gambler's Ruin

Simple random walk on $\{0, \ldots, N\}$ with absorbing barriers. The harmonic system $h(k) = \frac{1}{2}h(k-1) + \frac{1}{2}h(k+1)$, $h(0) = 0$, $h(N) = 1$ gives $h(k) = k/N$ for the probability of reaching N before 0. Optional stopping on $M_n = X_{n\wedge\tau}$ confirms: $k = N \cdot \mathbb{P}_k(\tau_N < \tau_0)$, so $\mathbb{P}_k(\tau_N < \tau_0) = k/N$.
For the biased walk with $\mathbb{P}(\text{step} + 1) = p \neq 1/2$, $M_n = (q/p)^{X_n}$ is a martingale and optional stopping gives $\mathbb{P}_k(\tau_N < \tau_0) = \frac{(q/p)^k - 1}{(q/p)^N - 1}$.

14.2 Azuma–Hoeffding and McDiarmid Inequalities

Theorem: Azuma–Hoeffding Inequality

Let $\{M_n\}$ be a martingale with $|M_k - M_{k-1}| \le c_k$ a.s. Then for all $t > 0$:

$$\mathbb{P}(M_n - M_0 \ge t) \le \exp\left(-\frac{t^2}{2\sum_{k=1}^{n} c_k^2}\right).$$

Proof. By the moment generating function approach: for each k, since $M_k - M_{k-1}$ is bounded in $[-c_k, c_k]$ and mean-zero given $\mathcal{F}_{k-1}$, Hoeffding's lemma gives $\mathbb{E}[e^{s(M_k - M_{k-1})} \mid \mathcal{F}_{k-1}] \le e^{s^2 c_k^2/2}$. Then e^{sM_n} is a submartingale after appropriate scaling: $\mathbb{E}[e^{sM_n}] \le e^{sM_0} \cdot e^{s^2 \sum c_k^2/2}$. Markov's inequality gives $\mathbb{P}(M_n - M_0 \ge t) \le e^{-st + s^2 \sum c_k^2/2}$. Optimizing over $s = t/\sum c_k^2$ gives the result. □

Theorem: McDiarmid's Inequality

Let $X_1, \dots, X_n$ be independent and f satisfy $\sup |f(\dots, x_k, \dots) - f(\dots, x_k', \dots)| \le c_k$. Then:

$$\mathbb{P}(f - \mathbb{E}[f] \ge t) \le \exp\left(-\frac{2t^2}{\sum_k c_k^2}\right).$$

14.3 Central Limit Theorem for Markov Chains

For an ergodic Markov chain and function f with $\pi(f) = 0$, the **CLT for Markov chains** states:

$$\frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} f(X_k) \xrightarrow{d} \mathcal{N}(0, \sigma_f^2),$$

where $\sigma_f^2 = \mathrm{Var}_\pi(f) + 2\sum_{n=1}^{\infty} \mathrm{Cov}_\pi(f(X_0), f(X_n))$ is the **asymptotic variance**. For reversible chains, $\sigma_f^2 = \langle f, (I + P)(I - P)^{-1} f \rangle_\pi$.

Chapter 15

Markov Chains in Statistical Physics

15.1 The Ising Model

Definition: Ising Model

Let $G = (V, E)$ be a finite graph. A configuration $\sigma \in \{-1, +1\}^V$ has Hamiltonian $H(\sigma) = -\sum_{\{i,j\}\in E} \sigma_i \sigma_j$ and Gibbs measure

$$\pi_\beta(\sigma) = \frac{e^{-\beta H(\sigma)}}{Z_\beta}.$$

At high temperature ($\beta \ll 1$) the measure is nearly uniform; at low temperature ($\beta \gg 1$) it concentrates on ground states. The critical temperature β_c marks the phase transition: on $\mathbb{Z}^2$, Onsager computed $\beta_c = \frac{1}{2}\ln(1+\sqrt{2})$.

Theorem: Mixing Time for Ising on Trees

For the Ising model on any tree with n vertices at any temperature, the Glauber dynamics mixes in $O(n \log n)$ steps.

Theorem: Slow Mixing Below Critical Temperature

For the Ising model on the $n \times n$ torus at $\beta > \beta_c$, the Glauber dynamics has mixing time $\exp(\Omega(n))$, due to metastability between the two ordered phases.

15.2 The Swendsen–Wang and Heat-Bath Algorithms

Glauber dynamics (heat-bath): at each step, choose $v \in V$ uniformly, resample σ_v from its conditional distribution given neighbors:

$$\mathbb{P}(\sigma_v = +1 \mid \sigma_{N(v)}) = \frac{e^{\beta \sum_{u\sim v} \sigma_u}}{e^{\beta \sum_{u\sim v} \sigma_u} + e^{-\beta \sum_{u\sim v} \sigma_u}}.$$

Swendsen–Wang: add each edge $\{i, j\}$ with $\sigma_i = \sigma_j$ independently with probability $1 - e^{-2\beta}$; identify clusters; assign each cluster a random spin. This algorithm mixes in $O(\log n)$ steps at all temperatures on the 2D Ising model.

Chapter 16

Harris Chains and General State Spaces

16.1 Minorization Conditions

Definition: Small Set and Minorization

Set $C \in \mathcal{B}$ satisfies the **minorization condition** if $\exists\, n_0, \varepsilon > 0$, probability measure ν:

$$P^{n_0}(x, \cdot) \geq \varepsilon\, \nu(\cdot) \qquad \forall x \in C.$$

Definition: Foster–Lyapunov Drift Condition

A chain satisfies the **drift condition** with Lyapunov function $V \geq 1$, constants $\lambda \in (0,1)$, $b < \infty$, and small set C:

$$PV(x) \leq \lambda V(x) + b\mathbf{1}_C(x).$$

Theorem: Geometric Ergodicity

If the minorization and drift conditions both hold, the chain is **geometrically ergodic**: $\exists\, r < 1$, $R : \mathcal{X} \to \mathbb{R}_+$ such that $\|P^n(x,\cdot) - \pi\|_{\mathrm{TV}} \leq R(x) r^n$. Moreover $R(x) = O(V(x))$.

Example: AR(1) Process

The AR(1) process $X_{n+1} = \rho X_n + \varepsilon_n$ with $|\rho| < 1$ and $\varepsilon_n \sim \mathcal{N}(0, \sigma^2)$ satisfies the drift condition with $V(x) = x^2 + 1$: $PV(x) = \rho^2 x^2 + \sigma^2 + 1 \leq \rho^2 V(x) + \sigma^2$. With $\lambda = \rho^2 < 1$ and $C = [-M, M]$ for large M, geometric ergodicity follows. The stationary distribution is $\mathcal{N}(0, \sigma^2/(1-\rho^2))$.

16.2 Regeneration Theory

The minorization condition yields a **regeneration construction**: at each visit to C, the chain independently starts fresh from ν with probability ε. These renewal epochs $\{\tau_k\}$ make successive tours i.i.d., enabling:

1. The **ergodic theorem**: $\frac{1}{n}\sum_{k=0}^{n-1} f(X_k) \to \pi(f)$ a.s.
2. The **CLT**: $\sqrt{n}(\bar{f}_n - \pi(f)) \xrightarrow{d} \mathcal{N}(0, \sigma_f^2)$.
3. **Variance estimation**: σ_f^2 can be consistently estimated from the batch means within regeneration cycles.

Chapter 17

Computational Algorithms

17.1 Pseudocode for Core Algorithms

Algorithm 1: Metropolis–Hastings MCMC

Input: Target π, proposal Q, initial state x_0, steps N
Output: Chain $\{x_0, \dots, x_N\}$

$x \leftarrow x_0$;
for $n = 1$ **to** N **do**
  Sample $y \sim Q(x, \cdot)$;
  $\alpha \leftarrow \min\left(1, \dfrac{\pi(y)Q(y,x)}{\pi(x)Q(x,y)}\right)$;
  Sample $u \sim \text{Uniform}(0,1)$;
  if $u \le \alpha$ **then**
    $x \leftarrow y$
  else
    $x \leftarrow x$
  Record $x_n \leftarrow x$;
return $\{x_0, \dots, x_N\}$;

Algorithm 2: Gibbs Sampler (Random Scan)

Input: Target π on $\mathcal{X}_1 \times \cdots \times \mathcal{X}_d$, initial $\mathbf{x}^{(0)}$, steps N
Output: Samples

$\mathbf{x} \leftarrow \mathbf{x}^{(0)}$;
for $n = 1$ **to** N **do**
  Sample $k \sim \text{Uniform}\{1, \dots, d\}$;
  Sample $x_k \sim \pi(x_k \mid x_{-k})$;
  $\mathbf{x}^{(n)} \leftarrow \mathbf{x}$ with updated k-th coordinate;
return $\{\mathbf{x}^{(0)}, \dots, \mathbf{x}^{(N)}\}$;

Algorithm 3: Propp–Wilson Coupling from the Past

Input: Monotone Markov chain on $\mathcal{X}$, initial $T = 1$
Output: Exact sample from π

repeat
  Generate random maps ϕ_t for $t = -T+1, \dots, 0$ (common randomness);
  for *each starting state* $x \in \mathcal{X}$ **do**
    Compute $X_0^{(x)} = \phi_0 \circ \cdots \circ \phi_{-T+1}(x)$;
  if *all* $X_0^{(x)}$ *agree on value* v **then**
    return v
  else
    $T \leftarrow 2T$
until *coalescence*;

17.2 Diagnostics and Convergence Assessment

In practice, mixing time cannot be computed directly; one uses empirical diagnostics:

1. **Autocorrelation time**: $\tau_{\text{int}} = 1 + 2\sum_{k=1}^{\infty} \rho(k)$ where $\rho(k) = \text{Corr}(f(X_0), f(X_k))$. Effective sample size $\text{ESS} = N/\tau_{\text{int}}$.

2. **Gelman–Rubin** $\hat{R}$: Run m chains from overdispersed starts. Compute between-chain variance B and within-chain variance W. Then $\hat{R} = \sqrt{((n-1)/n + B/(nW))} \approx 1$ at convergence.

3. **Trace plots**: Visual inspection of $\{f(X_n)\}$ for mixing and stationarity.

Remark: Burn-in and Warm Start

Chains are run with a burn-in of $B \approx t_{\text{mix}}$ steps discarded. A **warm start** initializing from a Laplace approximation reduces the required burn-in. The regenerative structure provides a principled definition: burn-in ends at the first regeneration time τ_1.

Chapter 18

Spectral Methods: Advanced Topics

18.1 Expander Mixing Lemma

The **expander mixing lemma** quantifies how well an expander graph approximates a random bipartite graph in terms of edge distribution.

Theorem: Expander Mixing Lemma

Let G be a d-regular graph on N vertices with second-largest eigenvalue λ_2 of the adjacency matrix. For any $S, T \subseteq V$:

$$\left| e(S,T) - \frac{d|S||T|}{N} \right| \leq \lambda_2 \sqrt{|S||T|},$$

where $e(S,T)$ is the number of edges between S and T.

Proof. Let $\mathbf{1}_S$ and $\mathbf{1}_T$ be indicator vectors. The adjacency matrix A has eigendecomposition $A = \sum_k \lambda_k \psi_k \psi_k^\top$ where $\{\psi_k\}$ are orthonormal eigenvectors. The first eigenvector is $\psi_1 = N^{-1/2}\mathbf{1}$ with $\lambda_1 = d$.

Then $e(S,T) = \mathbf{1}_S^\top A \mathbf{1}_T = \sum_k \lambda_k (\mathbf{1}_S \cdot \psi_k)(\mathbf{1}_T \cdot \psi_k)$. The $k = 1$ term gives $d|S||T|/N$. For the remaining terms, $|\sum_{k \geq 2} \lambda_k (\mathbf{1}_S \cdot \psi_k)(\mathbf{1}_T \cdot \psi_k)| \leq \lambda_2 \|\mathbf{1}_S\|_2 \|\mathbf{1}_T\|_2 = \lambda_2 \sqrt{|S||T|}$ by Cauchy–Schwarz. $\square$

18.2 Ramanujan Graphs

Definition: Ramanujan Graph

A d-regular graph G is a **Ramanujan graph** if every eigenvalue $\lambda \neq \pm d$ of its adjacency matrix satisfies $|\lambda| \leq 2\sqrt{d-1}$.

The Alon–Boppana theorem shows that $2\sqrt{d-1}$ is the infimum of λ_2 over all infinite families of d-regular graphs. Ramanujan graphs achieve this bound and therefore have the optimal spectral gap $\gamma = 1 - (2\sqrt{d-1})/d$ among d-regular graphs.

Lubotzky, Phillips, and Sarnak (1988) constructed explicit Ramanujan graphs using the arithmetic of quaternion algebras over $\mathbb{Q}$ and deep number-theoretic results (Ramanujan's conjecture, proved by Deligne). For p, q distinct primes with $p \equiv q \equiv 1 \pmod 4$, the LPS graph $X^{p,q}$ is a $(p+1)$-regular Ramanujan graph on $q(q^2-1)/2$ vertices.

18.3 Geometric Random Walks

For convex bodies $K \subseteq \mathbb{R}^d$, sampling from the uniform distribution on K is a fundamental problem in high-dimensional geometry and randomized algorithms.

Algorithm: Ball Walk

At state $x \in K$:

1. Sample $y \sim \text{Uniform}(B(x, \delta))$ where $B(x,\delta)$ is the ball of radius δ.
2. If $y \in K$, move to y; else stay at x.

Theorem: Mixing of Ball Walk (Lovász–Simonovits)

For a convex body $K \subseteq \mathbb{R}^d$ with $B(0,r) \subseteq K \subseteq B(0,R)$, the ball walk with $\delta = r/(10d)$ mixes in $O(d^2 (R/r)^2 \log(1/\varepsilon))$ steps from a warm start. With $R/r = O(1)$ (well-rounded body), the mixing time is $O(d^2 \log(d/\varepsilon))$.

The proof uses the isoperimetric inequality for convex bodies (KLS conjecture, now Klartag–Lehec 2022) and the connection between conductance and the spectral gap.

18.4 Langevin Monte Carlo

For smooth log-concave targets $\pi \propto e^{-U}$ on $\mathbb{R}^d$, the **Langevin algorithm** discretizes the Langevin SDE:

$$dX_t = -\nabla U(X_t)\,dt + \sqrt{2}\,dW_t.$$

The **unadjusted Langevin algorithm** (ULA) uses the Euler–Maruyama discretization:

$$X_{n+1} = X_n - \eta \nabla U(X_n) + \sqrt{2\eta}\,\xi_n, \qquad \xi_n \sim \mathcal{N}(0, I_d).$$

Theorem: Convergence of ULA (Dalalyan 2017)

For U m-strongly convex and M-smooth ($mI \preceq \nabla^2 U \preceq MI$), with step size $\eta \le 2/(m+M)$:

$$\|\mathcal{L}(X_n) - \pi\|_{\text{TV}}^2 \le (1 - \eta m)^n \cdot W_2^2(\mathcal{L}(X_0), \pi)/(2) + O(\eta d).$$

To achieve $\|\cdot - \pi\|_{\text{TV}} \le \varepsilon$, the number of steps is $O(\kappa d/\varepsilon^2)$ where $\kappa = M/m$ is the condition number.

Chapter 19

Stochastic Dominance and Monotone Chains

19.1 Stochastic Ordering on Lattices

When $(\mathcal{X}, \preceq)$ is a partially ordered set, we can define stochastic ordering and monotone Markov chains.

Definition: FKG Lattice Condition

A probability measure μ on a finite distributive lattice $(\mathcal{X}, \preceq)$ satisfies the **FKG condition** if for all $x, y \in \mathcal{X}$:

$$\mu(x)\mu(y) \leq \mu(x \wedge y)\mu(x \vee y),$$

where $x \wedge y$ and $x \vee y$ are the meet and join operations.

Theorem: FKG Inequality

If μ satisfies the FKG condition, then for any two increasing functions $f, g : \mathcal{X} \to \mathbb{R}$:

$$\mathbb{E}_\mu[fg] \geq \mathbb{E}_\mu[f]\,\mathbb{E}_\mu[g].$$

The FKG inequality is fundamental in statistical physics (Harris–FKG theorem for percolation and Ising models) and provides a powerful tool for correlation inequalities.

19.2 Monotone Coupling and Sandwiching

A Markov chain on $(\mathcal{X}, \preceq)$ is **monotone** if whenever $x \preceq y$, the corresponding coupled chains satisfy $X_n \preceq Y_n$ a.s. for all $n \geq 0$.

Theorem: Monotone CFTP

For a monotone Markov chain on a finite ordered set, Coupling from the Past terminates with probability 1 if and only if the chain run from the maximum and minimum elements of $\mathcal{X}$ eventually coalesce. The algorithm need only track two chains (from max and min) rather than $|\mathcal{X}|$.

This drastically reduces the computational cost of exact sampling for monotone chains. For the Ising model on a finite lattice at any temperature, the Glauber dynamics is monotone and CFTP can be implemented efficiently.

19.3 Interacting Particle Systems

Interacting particle systems are spatially extended Markov chains where particles at different sites evolve with local interactions.

Definition: Exclusion Process

The **simple exclusion process** on a graph $G = (V, E)$ with k particles: particles occupy k distinct vertices; at each step, choose a particle and an adjacent edge, and move the particle to the neighboring site if it is unoccupied. The stationary distribution is uniform over all $\binom{|V|}{k}$ configurations.

Theorem: Mixing Time for Exclusion (Wilson 2004)

The mixing time of the simple exclusion process with k particles on the complete graph K_n is $\Theta(n \log \min(k, n-k))$, exhibiting a cutoff at this time.

Chapter 20

Selected Open Problems and Research Directions

20.1 The KLS Conjecture

The **Kannan–Lovász–Simonovits (KLS) conjecture** (1995) states: there exists a universal constant $C > 0$ such that for any convex body $K \subseteq \mathbb{R}^d$ with centroid at the origin and $n \geq 2d$ points $x_1, \ldots, x_n$ drawn from the uniform distribution on K, the empirical distribution satisfies an isoperimetric inequality with constant $C \cdot d^{-1/2}$.

Equivalently, the KLS conjecture predicts that the Poincaré constant of the uniform distribution on any convex body is $O(d)$, implying that the mixing time of the ball walk is $O(d^2)$.

The conjecture was almost fully resolved by Klartag (2023), who showed the Poincaré constant is $O(d)$ up to logarithmic factors. This breakthrough settled a 30-year-old open problem and has implications for sampling algorithms in high dimensions.

20.2 Mixing and Computation

Many algorithmic problems in theoretical computer science reduce to understanding mixing times:

1. **Approximate counting**: Counting the number of independent sets, proper colorings, or perfect matchings in a graph. The correlation decay method and MCMC both play central roles.

2. **Ferromagnetic vs. antiferromagnetic models**: For ferromagnetic Ising/Potts, efficient MCMC algorithms exist at all temperatures. For antiferromagnetic models (hard-core gas, proper colorings), the picture depends on the activity/temperature and is connected to the uniqueness threshold on the infinite regular tree.

3. **Shattering threshold**: At low activities/temperatures, the Gibbs measure has a unique phase and MCMC mixes rapidly. Near and above the shattering thresh-

old, the measure fragments into exponentially many metastable states and MCMC provably fails.

20.3 Non-Reversible Chains and Acceleration

Reversibility is convenient for theory but not required for convergence. **Non-reversible Markov chains** can mix faster than any reversible chain.

Theorem: Acceleration by Non-Reversibility (Chen–Lovász–Pak)

For any reversible chain P with spectral gap γ, there exists a non-reversible chain $\tilde{P}$ with the same stationary distribution such that the absolute spectral gap satisfies $\tilde{\gamma}^* = O(\sqrt{\gamma})$, yielding a quadratic speedup $t_{\text{mix}}(\tilde{P}) = O(\sqrt{t_{\text{mix}}(P)})$.

The construction uses **lifting**: embedding the state space into a larger space with an auxiliary direction variable, creating a current that breaks detailed balance. The **zig-zag process** and **bouncy particle sampler** are continuous-time non-reversible samplers with provably faster mixing.

Chapter 21

Renewal Theory and Regenerative Processes

21.1 Renewal Sequences

Renewal theory studies sequences of independent, identically distributed waiting times between successive occurrences of an event, and provides a powerful lens for re-deriving and sharpening many Markov chain limit theorems.

Definition: Renewal Process

Let $\xi_1, \xi_2, \dots$ be i.i.d. nonnegative random variables with $\mathbb{P}(\xi_1 = 0) < 1$, representing inter-renewal times. Set $S_0 = 0$, $S_n = \xi_1 + \cdots + \xi_n$. The **renewal process** is $N(t) = \max\{n : S_n \le t\}$, the number of renewals by time t. The sequence is called a **renewal sequence** if we instead track the indicator process $u_n = \mathbb{P}(n \in \{S_0, S_1, S_2, \dots\})$.

For an irreducible, aperiodic, positive recurrent Markov chain (X_n) started at a fixed state i, the successive return times to i form the inter-renewal times of a renewal process, and $u_n = \mathbb{P}_i(X_n = i)$ is precisely the renewal sequence associated with the return-time distribution. This identification is the bridge between renewal theory and the convergence theory for Markov chains developed in earlier chapters.

Theorem: Renewal Equation

The renewal sequence satisfies

$$u_n = \sum_{k=1}^{n} f_k \, u_{n-k}, \qquad n \ge 1, \quad u_0 = 1,$$

where $f_k = \mathbb{P}(\xi_1 = k)$ is the inter-renewal distribution.

Proof. Condition on the time k of the first renewal after time 0: $\mathbb{P}(n \in \text{renewals} \mid \text{first renewal at } k) = u_{n-k}$ by the renewal property (independence and identical distribution of subsequent inter-renewal times), and summing over $k = 1, \dots, n$ with

weights f_k gives the stated identity. $\square$

21.2 The Renewal Theorem

Theorem: Discrete Renewal Theorem (Erdős–Feller–Pollard)

Suppose the inter-renewal distribution $(f_k)_{k\ge 1}$ is aperiodic (i.e., $\gcd\{k : f_k > 0\} = 1$) with mean $\mu = \sum_k k f_k \in (0, \infty]$. Then

$$\lim_{n\to\infty} u_n = \frac{1}{\mu},$$

with the convention $1/\infty = 0$.

This theorem is the renewal-theoretic restatement of the basic limit theorem for Markov chains: $\mathbb{P}_i(X_n = i) \to 1/\mathbb{E}_i[T_i^+]$, where T_i^+ is the first return time to i, recovering $\pi(i) = 1/\mathbb{E}_i[T_i^+]$ for the stationary distribution.

Definition: Key Renewal Theorem

For a function $g : \mathbb{N} \to \mathbb{R}$ that is directly Riemann integrable (in the discrete sense, summable with $\sum |g(k)| < \infty$), the **key renewal theorem** states

$$\sum_{k=0}^{n} g(n-k)\, u_k \;\longrightarrow\; \frac{1}{\mu}\sum_{k=0}^{\infty} g(k), \qquad n \to \infty.$$

The key renewal theorem is the workhorse for proving convergence rates: many quantities of interest in Markov chain theory (occupation probabilities, expected costs accrued between regenerations) can be written as a convolution against the renewal sequence, and the key renewal theorem immediately yields their limiting behavior.

21.3 Regenerative Processes and the Renewal-Reward Theorem

Definition: Regenerative Process

A stochastic process $(Y_t)_{t\ge 0}$ is **regenerative** with respect to a renewal sequence (S_n) if the post-S_n process $(Y_{S_n+t})_{t\ge 0}$ has the same law as $(Y_t)_{t\ge 0}$ and is independent of $(S_1, \dots, S_n)$, for every n.

Every irreducible positive recurrent Markov chain is regenerative with respect to the return times to any fixed reference state, which is precisely the mechanism underlying the renewal-theoretic proofs of the ergodic theorem in Chapter 4.

Theorem: Renewal-Reward Theorem

Let (Y_t) be regenerative with renewal epochs $S_0 = 0 < S_1 < S_2 < \cdots$, mean cycle length $\mu = \mathbb{E}[S_1]$, and let $R_n = \int_{S_{n-1}}^{S_n} h(Y_s)\,ds$ be the reward accumulated in cycle n for some bounded measurable h. If $\mathbb{E}[|R_1|] < \infty$, then almost surely

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t h(Y_s)\,ds = \frac{\mathbb{E}[R_1]}{\mu}.$$

Proof sketch. By the strong law of large numbers applied to the i.i.d. cycle rewards $R_1, R_2, \ldots$ and i.i.d. cycle lengths $\xi_1, \xi_2, \ldots = S_1 - S_0, S_2 - S_1, \ldots$, we have $\frac{1}{n}\sum_{k=1}^{n} R_k \to \mathbb{E}[R_1]$ and $\frac{1}{n}S_n \to \mu$ a.s. Sandwiching t between consecutive renewal times $S_{N(t)} \le t < S_{N(t)+1}$ and using $N(t)/t \to 1/\mu$ gives the result after a routine truncation argument controlling the boundary cycle. $\square$

This theorem unifies and extends the ergodic theorem for Markov chains: taking $Y_t = X_{\lfloor t \rfloor}$ and $h = \mathbf{1}_A$ recovers the long-run fraction of time spent in a set A, while taking h to be a cost function recovers long-run average cost formulas used throughout queueing and inventory theory.

21.4 Renewal Theory and Markov Chain Hitting Times

Example: Alternating Renewal Process and Availability

Consider a machine alternating between "up" periods of length $U_1, U_2, \ldots$ (i.i.d.) and "down" periods $D_1, D_2, \ldots$ (i.i.d., independent of the U_k). This is an alternating renewal process with cycle length $\xi_k = U_k + D_k$. By the renewal-reward theorem with $h(\text{state}) = \mathbf{1}\{\text{up}\}$, the long-run **availability** is

$$A = \frac{\mathbb{E}[U_1]}{\mathbb{E}[U_1] + \mathbb{E}[D_1]}.$$

This formula, ubiquitous in reliability engineering, is an instance of the renewal-reward theorem applied to a two-state regenerative process, and it coincides with the stationary probability of the "up" state when the alternating process is recast as a continuous-time Markov chain on $\{\text{up}, \text{down}\}$.

Proposition: Inspection Paradox

Let (S_n) be a renewal process with i.i.d. inter-renewal times of mean μ and variance σ^2, and fix a large time t. The length of the renewal interval straddling t converges in distribution, as $t \to \infty$, to the **size-biased** distribution with density $x f(x)/\mu$ (where f is the density of ξ_1), which has mean $\mu + \sigma^2/\mu \ge \mu$.

The inspection paradox explains a recurring counterintuitive phenomenon in applied Markov chain modeling: an observer arriving at a "random" time tends to land

in a longer-than-average renewal interval, because longer intervals occupy more of the time axis. This has direct consequences for waiting-time analysis in $M/G/1$-type queueing chains studied in Chapter 13.

21.5 Markov Renewal Processes

Definition: Markov Renewal Process

A **Markov renewal process** (or semi-Markov process) is a pair sequence $(X_n, T_n)_{n \geq 0}$ where (X_n) is a Markov chain on a countable state space with transition matrix P, and given the sequence (X_n), the sojourn times $T_{n+1} - T_n$ are independent with distribution $G_{X_n, X_{n+1}}$ depending only on the current and next state.

Semi-Markov processes generalize continuous-time Markov chains by allowing non-exponential sojourn times, at the cost of losing the memoryless property within each sojourn. The embedded jump chain (X_n) retains the full Markov chain machinery developed in this volume, while the sojourn-time structure is handled via renewal theory; this two-layer decomposition (Markov skeleton plus renewal-theoretic timing) is the standard technique for analyzing semi-Markov models in reliability theory and performance evaluation.

Chapter 22

Large Deviations for Markov Chains

22.1 The Large Deviation Principle

While the law of large numbers and central limit theorem (Chapter 15) describe typical and moderate fluctuations of additive functionals of a Markov chain, large deviation theory quantifies the exponentially small probability of atypical behavior.

Definition: Large Deviation Principle (LDP)

A sequence of probability measures (μ_n) on $\mathbb{R}^d$ satisfies a **large deviation principle** with rate function $I:\mathbb{R}^d\to[0,\infty]$ and speed n if I is lower semicontinuous and, for every Borel set A,

$$-\inf_{x\in A^\circ} I(x) \;\le\; \liminf_{n\to\infty}\frac{1}{n}\log \mu_n(A) \;\le\; \limsup_{n\to\infty}\frac{1}{n}\log \mu_n(A) \;\le\; -\inf_{x\in \bar{A}} I(x),$$

where A° and $\bar{A}$ are the interior and closure of A.

Definition: Logarithmic Moment Generating Function

For a Markov chain (X_n) on a finite state space with transition matrix P and a function $f:\mathcal{X}\to\mathbb{R}$, define the **tilted matrix** $P_\lambda(x,y)=P(x,y)e^{\lambda f(y)}$ and let

$$\Lambda(\lambda)=\log \rho(P_\lambda),$$

where $\rho(P_\lambda)$ is the Perron eigenvalue (spectral radius) of P_λ.

Theorem: Gärtner–Ellis / Donsker–Varadhan for Markov Additive Functionals

Let (X_n) be an irreducible Markov chain on a finite state space and $S_n=\sum_{k=1}^n f(X_k)$. Then (S_n/n) satisfies an LDP with speed n and rate function

$$I(x)=\sup_{\lambda\in\mathbb{R}}\{\lambda x-\Lambda(\lambda)\},$$

the Legendre–Fenchel transform of Λ. In particular, Λ is convex, smooth on a neighborhood of 0, and I is convex, vanishes at $x = \mathbb{E}_\pi[f(X)]$, and is strictly convex near that point.

Proof sketch. By the Perron–Frobenius theorem applied to the nonnegative irreducible matrix P_λ, $\rho(P_\lambda)$ is a simple eigenvalue with positive left and right eigenvectors, and $\mathbb{E}[e^{\lambda S_n}] \asymp \rho(P_\lambda)^n$ up to sub-exponential corrections coming from boundary terms involving the eigenvectors. This gives $\frac{1}{n}\log \mathbb{E}[e^{\lambda S_n}] \to \Lambda(\lambda)$, and the Gärtner–Ellis theorem (a generalization of Cramér's theorem to dependent sequences with well-behaved limiting log-moment generating function) then yields the LDP with rate function the Legendre transform of Λ. $\square$

22.2 The Donsker–Varadhan Variational Formula

An alternative, more probabilistic route to large deviations for Markov chains proceeds via relative entropy between probability measures on the state space.

Definition: Relative Entropy Rate Function

For an irreducible Markov chain on a finite state space with transition matrix P, define for any probability measure ν on $\mathcal{X}$:

$$H(\nu \mid P) = \sup_{u > 0} \sum_{x} \nu(x) \log \frac{u(x)}{(Pu)(x)},$$

where the supremum is over strictly positive functions $u : \mathcal{X} \to (0, \infty)$.

Theorem: Donsker–Varadhan LDP for Empirical Measures

Let $L_n = \frac{1}{n}\sum_{k=1}^{n} \delta_{X_k}$ be the empirical occupation measure of an irreducible Markov chain on a finite state space. Then (L_n) satisfies an LDP on the simplex of probability measures with speed n and rate function $I(\nu) = H(\nu \mid P)$. Moreover $H(\nu \mid P) = 0$ if and only if $\nu = \pi$, the stationary distribution.

This is the level-2 large deviation principle (large deviations of the empirical measure, as opposed to the level-1 principle for a single additive functional in the previous section); the level-1 rate function for f can be recovered from the level-2 rate function via the contraction principle, $I_f(x) = \inf\{H(\nu \mid P) : \mathbb{E}_\nu[f] = x\}$.

Remark: Connection to Relative Entropy and Reversibility

When P is reversible with respect to π, the rate function admits the more explicit Dirichlet-form representation

$$H(\nu \mid P) = \frac{1}{2}\sum_{x,y} \left(\sqrt{\nu(x)/\pi(x)} - \sqrt{\nu(y)/\pi(y)}\right)^2 \pi(x)P(x,y) + (\text{boundary terms}),$$

linking large deviation rate functions directly to the Dirichlet forms and spectral gap quantities of Chapter 5, and explaining why chains with larger spectral gap exhibit faster (in the LDP sense) concentration of the empirical measure around π.

22.3 Sanov's Theorem and Sample Path Large Deviations

Theorem: Sanov's Theorem (i.i.d. case)

If $X_1, X_2, \dots$ are i.i.d. with law μ on a finite alphabet, the empirical measures $L_n = \frac{1}{n}\sum_{k=1}^{n} \delta_{X_k}$ satisfy an LDP with rate function $I(\nu) = D_{\mathrm{KL}}(\nu \,\|\, \mu) = \sum_x \nu(x) \log(\nu(x)/\mu(x))$.

Sanov's theorem is the $P =$ i.i.d. special case of the Donsker–Varadhan principle: when $P(x, y) = \mu(y)$ for all x (independent of the current state), the relative entropy rate function $H(\nu \mid P)$ reduces exactly to the Kullback–Leibler divergence $D_{\mathrm{KL}}(\nu\|\mu)$, recovering the classical result as a degenerate case of the Markov chain theory.

Example: Large Deviations for a Two-State Chain

For the two-state chain on $\{0, 1\}$ with $P(0, 1) = p$, $P(1, 0) = q$, the empirical fraction of time in state 1, $\bar{X}_n = \frac{1}{n}\sum_{k=1}^{n} X_k$, satisfies an LDP with rate function

$$I(x) = x \log \frac{x}{\pi(1)} + (1 - x) \log \frac{1 - x}{\pi(0)} + (\text{correction from transition structure}),$$

which to leading order matches the i.i.d. Bernoulli$(\pi(1))$ rate function but includes corrections reflecting the positive correlation introduced by the Markov dependence; these corrections vanish in the weak-dependence limit $p, q \to 1/2$ where consecutive states become asymptotically independent.

22.4 Applications: Exit Times and Metastability

Large deviation theory provides sharp asymptotics for the time it takes a Markov chain to exit a domain or transition between metastable states, a problem central to the analysis of MCMC algorithms (Chapter 10) on multimodal target distributions.

Theorem: Freidlin–Wentzell-Type Exit Time Asymptotics

Let (X_n) be a Markov chain depending on a small parameter ε (e.g., a low-temperature Glauber dynamics with inverse temperature $\beta = 1/\varepsilon$) with a unique stable equilibrium x^* inside a domain D. Define the quasi-potential $V(x) =$ inf over paths from x^* to x of the large-deviation action functional. Then

the exit time τ_D from D satisfies

$$\lim_{\varepsilon \to 0} \varepsilon \log \mathbb{E}[\tau_D] = \min_{x \in \partial D} V(x) =: \bar{V}.$$

This result formalizes the exponential-in-$1/\varepsilon$ (equivalently, exponential-in-β) time required for low-temperature MCMC chains to escape a local mode and explains the practical phenomenon, observed already in Chapter 14's cutoff discussion, that mixing times for multimodal targets can be astronomically larger than the spectral gap of any single mode alone would suggest, because the bottleneck is governed by the large-deviation cost of crossing the energy barrier rather than by local mixing within a mode.

Chapter 23

Markov Chains on Infinite and Random Graphs

23.1 Recurrence and Transience Revisited: The Role of Geometry

Chapter 3 established recurrence and transience criteria for general countable-state Markov chains via the existence of excessive functions. We now specialize to simple random walk on graphs, where geometric and combinatorial structure of the underlying graph directly determines recurrence, transience, and the rate of escape.

Theorem: Pólya's Recurrence Theorem

Simple random walk on $\mathbb{Z}^d$ is recurrent for $d \in \{1, 2\}$ and transient for $d \geq 3$.

Proof sketch via generating functions. By the local central limit theorem, $\mathbb{P}(X_{2n} = 0) \sim Cn^{-d/2}$ for simple random walk on $\mathbb{Z}^d$. Recurrence is equivalent to $\sum_n \mathbb{P}(X_{2n} = 0) = \infty$ (Chapter 3's recurrence criterion via expected number of returns), which holds iff $d/2 \leq 1$, i.e., $d \leq 2$. □

Definition: Effective Resistance

View a graph $G = (V, E)$ as an electrical network with unit resistance on each edge. The **effective resistance** $R_{\text{eff}}(a, b)$ between vertices a, b is the voltage difference required to drive a unit current from a to b.

Theorem: Random Walk–Electrical Network Correspondence

Simple random walk on a graph G started at a is transient to infinity, restricted to the event of never returning to a, with escape probability related to effective resistance by

$$\mathbb{P}_a(\text{never return to } a) = \frac{1}{d(a) R_{\text{eff}}(a, \infty)},$$

where $d(a)$ is the degree of a and $R_{\text{eff}}(a,\infty) = \lim_{n\to\infty} R_{\text{eff}}(a, \partial B_n)$ for a sequence of balls B_n exhausting G (interpreted as ∞ when the limiting resistance diverges, in which case the walk is recurrent).

This electrical network correspondence, due to Doyle and Snell, recasts the recurrence/transience dichotomy of Chapter 3 in terms of a purely deterministic combinatorial quantity (effective resistance), and provides a second, geometric proof of Pólya's theorem via direct resistance estimates on $\mathbb{Z}^d$.

23.2 Random Walk on Random Graphs: The Erdős–Rényi Model

Definition: Erdős–Rényi Random Graph

The **Erdős–Rényi random graph** $G(n,p)$ has vertex set $\{1,\dots,n\}$ and each of the $\binom{n}{2}$ possible edges present independently with probability p.

Theorem: Mixing Time of Random Walk on $G(n,p)$

For $p = p(n)$ satisfying $np \geq (1+\delta)\log n$ for some $\delta > 0$ (above the connectivity threshold), with high probability the simple random walk on $G(n,p)$ has mixing time

$$t_{\text{mix}} = O\left(\frac{\log n}{\log(np)}\right),$$

exhibiting the small-world phenomenon: mixing occurs almost as fast as on the complete graph.

The proof combines a spectral gap estimate for the (random) adjacency matrix of $G(n,p)$, obtained via concentration-of-measure arguments showing the random graph's normalized adjacency matrix is close in operator norm to its expectation, with the conductance/Cheeger machinery of Chapter 5; the random graph behaves, spectrally, like an expander once it is sufficiently dense, connecting this section directly to the Ramanujan graph and expander mixing material developed earlier in this volume.

23.3 Random Walk on Galton–Watson Trees

Definition: Galton–Watson Tree

A **Galton–Watson tree** with offspring distribution $(p_k)_{k\geq 0}$ is the family tree of a branching process: the root has a random number of children with law (p_k), and each subsequent individual independently has children according to the same law.

Theorem: Recurrence of Random Walk on GW Trees (Lyons)

Simple random walk on a Galton–Watson tree with offspring mean $m = \sum_k k p_k$, conditioned on non-extinction, is transient almost surely whenever $m > 1$ (the supercritical regime). When $m \le 1$, the tree is almost surely finite (or a single infinite path in the critical case conditioned to survive), and the walk is trivially recurrent or escapes to the unique boundary point.

The transience criterion for supercritical Galton–Watson trees is most cleanly proved via the Nash-Williams criterion: the tree, conditioned to survive, has at generation n roughly m^n vertices, giving exponentially diverging effective resistance lower bounds that, combined with the electrical network correspondence above, immediately yield transience for $m > 1$. This example showcases the power of combining branching process theory (a recurring theme across this textbook series) with the electrical/Markov chain correspondence developed in this chapter.

23.4 The Configuration Model and Local Weak Convergence

Definition: Configuration Model

Given a degree sequence $d_1, \dots, d_n$, the **configuration model** constructs a random (multi-)graph by attaching d_i "half-edges" to vertex i and pairing half-edges uniformly at random into edges.

Definition: Local Weak Convergence

A sequence of (random) graphs G_n rooted at a uniformly random vertex **converges locally weakly** to a rooted random graph (G_∞, o) if, for every r, the distribution of the r-ball around the root converges in distribution (as rooted graphs) to the r-ball around o in G_∞.

Theorem: Local Limit of the Configuration Model

If the empirical degree distribution of the configuration model converges to a limiting law (p_k) with finite mean, the configuration model converges locally weakly to the **unimodular Galton–Watson tree** with root offspring distribution (p_k) and non-root offspring distribution given by the size-biased law $\hat{p}_k = (k+1)p_{k+1}/\sum_j j p_j$.

Local weak convergence provides the rigorous mechanism by which random walk asymptotics on sparse random graphs (mixing time, spectral measure, return probabilities) can be deduced from the corresponding quantities on the limiting Galton–Watson tree, which is typically far more tractable; this technique unifies the Erdős–Rényi and configuration-model examples above as special cases of a single general principle, and connects back to the spectral methods of Chapter 5 through the method

of moments for the limiting spectral distribution of the adjacency matrix (the Kesten–McKay law for regular trees and its generalizations for irregular offspring distributions).

Remark: Outlook

The interplay between Markov chains and random graph geometry developed in this chapter is a rapidly active research area, connecting the foundational material of this volume to contemporary topics including community detection via random walks, graph neural network expressivity (where random walk return probabilities serve as graph invariants), and the statistical physics of disordered systems on sparse random graphs.

Chapter 24

Hidden Markov Models and Filtering

24.1 The Hidden Markov Model Framework

In many applications the Markov chain itself is not directly observed; instead, one observes a noisy or partial function of the chain at each time step, and the inferential task is to recover information about the hidden state from the observed data.

Definition: Hidden Markov Model

A **hidden Markov model (HMM)** consists of a hidden Markov chain $(X_n)_{n\ge 0}$ on a state space $\mathcal{X}$ with transition matrix P and initial distribution μ_0, together with observations $(Y_n)_{n\ge 0}$ taking values in an observation space $\mathcal{Y}$, such that conditionally on (X_n), the Y_n are independent with $Y_n \sim g(\cdot \mid X_n)$ for an **emission kernel** g.

The defining conditional independence structure, $Y_n \perp (X_m, Y_m)_{m\neq n} \mid X_n$, makes the joint process (X_n, Y_n) itself a Markov chain on the product space $\mathcal{X}\times\mathcal{Y}$, allowing all of the machinery of this volume (stationary distributions, ergodic theorems, coupling) to be brought to bear, while the central inferential problems concern the conditional law of X_n given the observed history $Y_{0:n} = (Y_0, \dots, Y_n)$.

24.2 The Forward Algorithm and Filtering Recursion

Definition: Filtering Distribution

The **filtering distribution** at time n is $\pi_n(x) = \mathbb{P}(X_n = x \mid Y_{0:n})$, the posterior law of the hidden state given all observations up to the present time.

Theorem: Filtering Recursion

The filtering distribution satisfies the recursive update

$$\pi_{n+1}(x) = \frac{g(Y_{n+1} \mid x) \sum_{x'} P(x', x)\, \pi_n(x')}{\sum_{x''} g(Y_{n+1} \mid x'') \sum_{x'} P(x', x'')\, \pi_n(x')},$$

consisting of a **prediction step** (propagating π_n forward via P) followed by a **Bayesian update step** (reweighting by the likelihood of the new observation).

Proof. By Bayes' rule and the conditional independence structure of the HMM,

$$\mathbb{P}(X_{n+1} = x \mid Y_{0:n+1}) \propto \mathbb{P}(Y_{n+1} \mid X_{n+1} = x) \cdot \mathbb{P}(X_{n+1} = x \mid Y_{0:n}),$$

and the Chapman–Kolmogorov-type identity $\mathbb{P}(X_{n+1} = x \mid Y_{0:n}) = \sum_{x'} P(x', x)\pi_n(x')$ follows from the Markov property of (X_n) together with the conditional independence of $Y_{0:n}$ and X_{n+1} given X_n. Normalizing gives the stated formula. □

This recursion, known in the engineering literature as the **forward algorithm** when the state space is finite, is a direct descendant of the Chapman–Kolmogorov equations of Chapter 2, modified to incorporate the likelihood reweighting from each new observation; its computational cost per step is $O(|\mathcal{X}|^2)$, matching one application of the transition matrix.

24.3 Smoothing and the Forward–Backward Algorithm

Definition: Smoothing Distribution

The **smoothing distribution** at time $k \le n$ is $\mathbb{P}(X_k = x \mid Y_{0:n})$, the posterior law of a past hidden state given the entire observed sequence up to a later time n.

Theorem: Forward–Backward Decomposition

Define the forward variables $\alpha_k(x) = \mathbb{P}(X_k = x, Y_{0:k})$ and backward variables $\beta_k(x) = \mathbb{P}(Y_{k+1:n} \mid X_k = x)$. Then the smoothing distribution factorizes as

$$\mathbb{P}(X_k = x \mid Y_{0:n}) = \frac{\alpha_k(x)\, \beta_k(x)}{\sum_{x'} \alpha_k(x')\, \beta_k(x')},$$

where α_k satisfies the forward recursion of the previous section (up to normalization) and β_k satisfies the backward recursion $\beta_k(x) = \sum_{x'} P(x, x')\, g(Y_{k+1} \mid x')\, \beta_{k+1}(x')$ with $\beta_n \equiv 1$.

The backward recursion runs the transition matrix machinery of Chapter 2 in reverse time, and the forward–backward algorithm's $O(n|\mathcal{X}|^2)$ total cost makes exact smoothing tractable for moderate state spaces, in sharp contrast to the exponential cost of naively summing over all hidden trajectories.

24.4 The Viterbi Algorithm: Most Likely State Sequence

Algorithm: Viterbi Algorithm

To find $\hat{x}_{0:n} = \arg\max_{x_{0:n}} \mathbb{P}(X_{0:n} = x_{0:n} \mid Y_{0:n})$, define $\delta_k(x) = \max_{x_{0:k-1}} \mathbb{P}(X_{0:k-1} = x_{0:k-1}, X_k = x, Y_{0:k})$ and recurse:

$$\delta_{k+1}(x) = g(Y_{k+1} \mid x) \cdot \max_{x'} \big[P(x', x)\, \delta_k(x')\big],$$

recording the maximizing x' at each step (backpointers), then trace back from $\arg\max_x \delta_n(x)$.

The Viterbi algorithm replaces the sum over predecessor states in the forward algorithm with a maximum, turning the filtering recursion into a dynamic program over the trellis of state sequences; this max-product (as opposed to sum-product) structure is an instance of the general max-plus algebra reformulation of shortest-path/longest-path problems, and the algorithm's correctness follows from the same optimality-of-subpaths argument used in shortest-path dynamic programming.

24.5 Parameter Estimation: The Baum–Welch Algorithm

When the transition matrix P and emission kernel g are unknown and must be estimated from observed data alone, the **Baum–Welch algorithm** (a special case of the Expectation–Maximization algorithm) provides an iterative maximum-likelihood procedure.

Theorem: EM Monotonicity for HMMs

Given current parameter estimates $\theta^{(t)} = (P^{(t)}, g^{(t)})$, define updated estimates $\theta^{(t+1)}$ by maximizing the expected complete-data log-likelihood

$$Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\theta^{(t)}}\big[\log \mathbb{P}_\theta(X_{0:n}, Y_{0:n}) \mid Y_{0:n}\big]$$

over θ. Then the observed-data log-likelihood is non-decreasing: $\log \mathbb{P}_{\theta^{(t+1)}}(Y_{0:n}) \geq \log \mathbb{P}_{\theta^{(t)}}(Y_{0:n})$.

Proof sketch. By Jensen's inequality applied to the concavity of log,

$$\log \mathbb{P}_\theta(Y_{0:n}) - \log \mathbb{P}_{\theta^{(t)}}(Y_{0:n}) \geq Q(\theta \mid \theta^{(t)}) - Q(\theta^{(t)} \mid \theta^{(t)})$$

for all θ (this is the standard EM lower-bound argument), and since $\theta^{(t+1)}$ maximizes $Q(\cdot \mid \theta^{(t)})$, the right side is nonnegative at $\theta = \theta^{(t+1)}$. □

For HMMs, the maximization step has closed form: the updated transition matrix entries $P^{(t+1)}(x, x')$ are proportional to the expected number of transitions from x to x' under the smoothing distribution computed via forward–backward, directly connecting the Baum–Welch update to the empirical transition count estimators of Chapter 2

used for fully observed chains, now weighted by posterior state probabilities rather than observed indicators.

24.6 Particle Filtering for Continuous and Large State Spaces

When $\mathcal{X}$ is continuous or too large for exact forward–backward computation, the filtering recursion is approximated using sequential Monte Carlo.

Algorithm: Bootstrap Particle Filter

Maintain N weighted particles $(x_n^{(i)}, w_n^{(i)})_{i=1}^N$ approximating π_n. At each step:

1. **Propagate**: sample $\tilde{x}_{n+1}^{(i)} \sim P(x_n^{(i)}, \cdot)$ for each i.
2. **Weight**: set $w_{n+1}^{(i)} \propto g(Y_{n+1} \mid \tilde{x}_{n+1}^{(i)})$.
3. **Resample**: draw $x_{n+1}^{(i)}$ from $\{\tilde{x}_{n+1}^{(j)}\}$ with probability proportional to $w_{n+1}^{(j)}$, then reset weights to $1/N$.

Theorem: Convergence of the Particle Filter

Under mild regularity conditions on g (bounded above and away from zero), the particle filter approximation $\hat{\pi}_n^N = \sum_i w_n^{(i)} \delta_{x_n^{(i)}}$ satisfies, for bounded test functions φ,

$$\mathbb{E}\left[\left(\hat{\pi}_n^N(\varphi) - \pi_n(\varphi)\right)^2\right] \leq \frac{C_n \|\varphi\|_\infty^2}{N}$$

for a constant C_n that, under additional mixing conditions on the hidden chain (a sufficiently large spectral gap, in the sense of Chapter 5), can be bounded uniformly in n, yielding particle filters whose accuracy does not degrade over long time horizons.

The resampling step is precisely a multinomial coupling of the weighted particle population, and the uniform-in-time stability result relies on a contraction argument for the resampling operation that closely parallels the coupling contraction arguments used to prove mixing time bounds in Chapters 6–7; intuitively, a chain that itself mixes rapidly prevents the particle approximation error from compounding over time, since the filter's own "forgetting" of the distant past mirrors the hidden chain's forgetting of its initial condition.

Remark: Outlook and Connections

Hidden Markov models and particle filtering sit at the confluence of nearly every technique developed in this volume: the Chapman–Kolmogorov equations and transition matrix algebra of Chapters 2–3, the coupling techniques of Chapters 6–7 (now applied to the resampling step rather than to the chain itself), and the spectral gap quantities of Chapter 5 (governing the long-time stability of the filter). This makes HMM filtering theory a natural capstone, illustrating how

the foundational tools of Markov chain theory extend beyond the autonomous chains studied in earlier chapters to the partially observed setting that dominates modern applications in signal processing, computational biology, and machine learning.

Chapter 25

Markov Decision Processes and Optimal Control

25.1 The Markov Decision Process Framework

We now introduce a controlled extension of the Markov chain model in which an agent influences the transition dynamics by selecting actions, with the goal of optimizing a long-run reward criterion. This framework subsumes the autonomous chains of earlier chapters as the degenerate case of a single available action.

Definition: Markov Decision Process

A **Markov decision process (MDP)** is a tuple $(\mathcal{X}, \mathcal{A}, P, r)$ where $\mathcal{X}$ is the state space, $\mathcal{A}$ is the action space, $P(x' \mid x, a)$ is the transition kernel giving the law of the next state given current state x and action a, and $r(x, a)$ is the one-step reward.

Definition: Policy and Induced Markov Chain

A (stationary, deterministic) **policy** is a map $\pi : \mathcal{X} \to \mathcal{A}$. Given a policy π, the process (X_n) evolves as the Markov chain with transition matrix $P^{\pi}(x, x') = P(x' \mid x, \pi(x))$, to which the entire apparatus of this volume (stationary distributions, mixing times, ergodic theorems) applies directly.

25.2 Discounted Reward Criterion and the Bellman Equation

Definition: Discounted Value Function

For a discount factor $\gamma \in (0, 1)$ and policy π, the **value function** is

$$V^{\pi}(x) = \mathbb{E}_x\left[\sum_{n=0}^{\infty} \gamma^n r(X_n, \pi(X_n))\right].$$

The **optimal value function** is $V^*(x) = \sup_\pi V^\pi(x)$.

Theorem: Bellman Optimality Equation

The optimal value function is the unique bounded solution of

$$V^*(x) = \max_{a \in \mathcal{A}} \left\{ r(x,a) + \gamma \sum_{x'} P(x' \mid x, a)\, V^*(x') \right\},$$

and any policy π^* achieving the maximum for every x is optimal: $V^{\pi^*} = V^*$.

Proof sketch. Define the **Bellman operator** $(\mathcal{T}V)(x) = \max_a \{r(x,a) + \gamma \sum_{x'} P(x' \mid x,a) V(x')\}$ on the space of bounded functions with sup norm. One checks $\mathcal{T}$ is a γ-contraction: $\|\mathcal{T}V_1 - \mathcal{T}V_2\|_\infty \le \gamma \|V_1 - V_2\|_\infty$, since the inner expressions for any fixed a differ by at most $\gamma\|V_1 - V_2\|_\infty$ in absolute value, and taking a max over a preserves this Lipschitz bound. By the Banach fixed point theorem, $\mathcal{T}$ has a unique fixed point, and a standard verification argument shows this fixed point equals V^*. $\square$

The contraction argument here is a direct descendant of the contraction mapping ideas used in Chapter 6 (coupling) and Chapter 9 (semigroup generators): just as γ-discounting guarantees a unique fixed point for the Bellman operator, the spectral gap of P^π guarantees geometric convergence to the stationary distribution for any fixed policy π.

25.3 Value Iteration and Policy Iteration

Algorithm: Value Iteration

Initialize $V_0 \equiv 0$. Iterate $V_{n+1} = \mathcal{T} V_n$ until convergence. By the contraction property, $\|V_n - V^*\|_\infty \le \gamma^n \|V_0 - V^*\|_\infty \to 0$ geometrically fast.

Algorithm: Policy Iteration

Initialize an arbitrary policy π_0. Iterate:

1. **Policy evaluation**: solve $V^{\pi_n}(x) = r(x, \pi_n(x)) + \gamma \sum_{x'} P(x' \mid x, \pi_n(x)) V^{\pi_n}(x')$ exactly (a linear system).
2. **Policy improvement**: set $\pi_{n+1}(x) = \arg\max_a \{r(x,a) + \gamma \sum_{x'} P(x' \mid x,a) V^{\pi_n}(x')\}$.

Theorem: Finite Convergence of Policy Iteration

For a finite state and action space, policy iteration produces a non-decreasing sequence of value functions $V^{\pi_0} \le V^{\pi_1} \le \cdots$, and since there are only finitely

many deterministic policies, the algorithm converges to π^* in finitely many iterations.

The policy evaluation step solves exactly the type of linear system encountered when computing expected hitting times and occupation measures in Chapter 3, with $(I - \gamma P^{\pi_n})$ playing the role of a discounted analogue of the fundamental matrix; the well-posedness of this linear system follows from $\gamma < 1$ ensuring $(I - \gamma P^{\pi_n})$ is invertible, since P^{π_n} is a stochastic matrix with spectral radius 1.

25.4 Average-Reward MDPs and the Connection to Stationary Distributions

Definition: Long-Run Average Reward

For a policy π inducing an irreducible, positive recurrent chain P^π with stationary distribution π^π, the **long-run average reward** is

$$g^\pi = \sum_x \pi^\pi(x)\, r(x, \pi(x)) = \mathbb{E}_{\pi^\pi}[r(X, \pi(X))].$$

Theorem: Average-Reward Optimality (Bellman Equation)

The optimal average reward $g^* = \max_\pi g^\pi$ and an associated **bias function** h^* satisfy

$$g^* + h^*(x) = \max_a \left\{ r(x,a) + \sum_{x'} P(x' \mid x, a)\, h^*(x') \right\},$$

and the maximizing policy is average-reward optimal.

This average-reward formulation is obtained from the discounted criterion in the limit $\gamma \to 1$, and g^π is exactly the long-run average reward guaranteed to exist by the ergodic theorem of Chapter 4 applied to the reward process $r(X_n, \pi(X_n))$; the bias function h^* plays a role directly analogous to the fundamental matrix deviation terms that appear in the central limit theorem for Markov chains (Chapter 15), both quantifying transient deviations from the long-run average.

25.5 Reinforcement Learning: Learning Without a Known Model

When the transition kernel P and reward function r are unknown and must be learned from interaction, the MDP framework gives rise to reinforcement learning algorithms.

Algorithm: Q-Learning

Maintain estimates $Q(x,a)$ of the optimal action-value function $Q^*(x,a) = r(x,a) + \gamma \sum_{x'} P(x' \mid x,a) \max_{a'} Q^*(x',a')$. Upon observing a transition (x,a,r,x'), update

$$Q(x,a) \leftarrow Q(x,a) + \alpha_n(x,a)\Big[r + \gamma \max_{a'} Q(x',a') - Q(x,a)\Big].$$

Theorem: Convergence of Q-Learning (Watkins–Dayan / Stochastic Approximation)

If every state-action pair is visited infinitely often and the step sizes satisfy the Robbins–Monro conditions $\sum_n \alpha_n(x,a) = \infty$ and $\sum_n \alpha_n(x,a)^2 < \infty$, then $Q(x,a) \to Q^*(x,a)$ almost surely.

Proof sketch. Q-learning is a stochastic approximation scheme for the fixed-point equation of the Bellman operator $\mathcal{T}$, which is a γ-contraction in sup norm as shown above. The Robbins–Monro step-size conditions are precisely those required for the Robbins–Siegmund almost-supermartingale convergence theorem to guarantee convergence of stochastic approximation iterates to the fixed point of a contraction, in the presence of bounded-variance noise from the random transitions and rewards. □

The "infinitely often visited" requirement is an ergodicity-type condition closely related to the recurrence criteria of Chapter 3: under an exploration policy inducing an irreducible chain on the state-action space, positive recurrence guarantees the required infinite visitation, linking the convergence theory of reinforcement learning algorithms directly back to the recurrence and ergodicity foundations laid in the early chapters of this volume.

Remark: Outlook

25.6 A Worked Example: The Optimal Stopping Problem

We close with a classical MDP that connects directly to the martingale theory of Chapter 16.

Example: Optimal Stopping as a Two-Action MDP

Consider a Markov chain (X_n) on a finite state space with a reward $h(x)$ for stopping in state x, and the option at each step to either stop (action `stop`, terminal reward $h(X_n)$) or continue (action `continue`, transitioning according to P with no immediate reward). This is an MDP with action space $\mathcal{A} = \{\texttt{stop}, \texttt{continue}\}$, and the Bellman equation specializes to

$$V^*(x) = \max\Big\{h(x),\ \mathbb{E}_x[V^*(X_1)]\Big\}.$$

The optimal policy stops the first time $V^*(X_n) = h(X_n)$, i.e., the first time continuing offers no further advantage.

Proposition: V^* is the Smallest Superharmonic Majorant

The value function V^* of the optimal stopping problem is the smallest function V satisfying simultaneously $V(x) \geq h(x)$ for all x and $V(x) \geq \mathbb{E}_x[V(X_1)]$ for all x (i.e., V is superharmonic for P and dominates h).

This characterization is the discrete-time, discrete-state analogue of the classical Snell envelope construction in continuous-time optimal stopping theory, and the superharmonicity condition $V \geq PV$ is exactly the excessive-function condition used in Chapter 3 to characterize transient and recurrent behavior; optimal stopping theory can thus be viewed as asking which excessive functions dominating h are tightest, with the stopped process $(X_{n\wedge\tau^*})$, where τ^* is the optimal stopping time, forming a martingale under the optimally stopped dynamics by the optional stopping theorem of Chapter 16.

Example: The House-Selling Problem

A seller receives i.i.d. offers $X_1, X_2, \ldots$ for a house and must decide, upon each offer, whether to accept (ending the process with reward X_n minus accumulated holding cost cn) or reject and wait for the next offer. The Bellman equation for the value-to-go $V(x)$ given a current offer x reduces to a threshold rule: accept iff $x \geq x^*$, where the threshold x^* solves $x^* = \mathbb{E}[\max(X, x^*)] - c$, balancing the holding cost c against the expected gain from waiting for a better offer. This threshold structure, a hallmark of monotone optimal stopping problems, generalizes directly to many sequential search and inventory models built on the same MDP scaffolding.

Markov decision processes and reinforcement learning represent the natural extension of Markov chain theory to settings with control and learning, and the mathematical machinery developed throughout this volume, contraction mappings, spectral gaps, ergodic theorems, and stochastic approximation, reappears throughout in service of optimality and convergence guarantees. Readers interested in the deeper mathematics of policy gradient methods, function ap-

proximation, and the convergence theory of modern deep reinforcement learning algorithms are referred to the companion volume on the mathematics of neural networks.

Concluding Remarks

This volume has developed the theory of Markov chains from first principles to modern quantitative results. The central thread connecting all parts is the interplay between three fundamental perspectives: the algebraic perspective (transition matrices, eigenvalues, stationary vectors), the probabilistic perspective (coupling, stopping times, hitting times), and the geometric perspective (conductance, expanders, Dirichlet forms).

The spectral gap γ emerges as the universal quantity governing convergence: it appears in the exponential decay of total variation distance, in the Poincaré inequality, in the conductance via Cheeger's inequality, and in the mixing time bounds from both the spectral and coupling approaches.

The later chapters extended this core machinery in three directions that mirror the trajectory of the field as a whole. First, renewal theory and regenerative processes (Chapter 19) recast the ergodic theorem as an instance of a more general renewal-reward principle, revealing the return-time structure underlying positive recurrence as the source of every long-run average computed in this volume. Second, large deviation theory (Chapter 20) sharpened the law of large numbers into exponential-rate statements governing rare events, exit times, and metastability, with the rate function's Legendre-transform structure and its Donsker–Varadhan relative-entropy form both tracing back to the spectral and Dirichlet-form quantities of Chapter 5. Third, the study of Markov chains on infinite and random graphs (Chapter 21) showed how the electrical network correspondence and local weak convergence translate purely combinatorial graph properties, expander mixing, effective resistance, branching structure, into sharp recurrence and mixing time dichotomies.

The final two chapters turned from autonomous evolution to inference and control. Hidden Markov models and particle filtering (Chapter 22) addressed the partially observed setting in which the chain itself must be inferred from noisy data, while Markov decision processes and reinforcement learning (Chapter 23) introduced an agent that actively shapes the transition dynamics through a chosen policy, reducing the control problem to a fixed-point equation solved by the same contraction-mapping principle that underlies coupling-based mixing time bounds. Both directions illustrate that the Markov property, conditional independence of the future from the past given the present, remains the organizing principle even once the chain is observed only indirectly or steered by external decisions.

Volume II takes these foundations into stochastic calculus, where Markov chains become diffusion processes, the generator becomes a differential operator, and the spectral gap theory extends to functional inequalities.

Appendix A

Linear Algebra and Perron–Frobenius Theory

A.1 Eigenvalues and Eigenvectors of Stochastic Matrices

This appendix collects the linear-algebraic facts about nonnegative and stochastic matrices used throughout the volume, in particular in Chapters 5, 9, and 10.

Definition: Nonnegative and Irreducible Matrices

A matrix $A \in \mathbb{R}^{N \times N}$ is **nonnegative** if $A_{ij} \geq 0$ for all i, j, and **irreducible** if for every i, j there exists n with $(A^n)_{ij} > 0$.

Theorem: Perron–Frobenius

Let A be a nonnegative irreducible $N \times N$ matrix. Then:

1. A has a real eigenvalue $\rho(A) > 0$ (the **Perron eigenvalue**) equal to its spectral radius;
2. $\rho(A)$ is simple (algebraic multiplicity one) and has a strictly positive eigenvector, unique up to scaling;
3. every other eigenvalue λ satisfies $|\lambda| \leq \rho(A)$, with equality possible only if A is periodic.

For a stochastic matrix P (row sums equal to 1), $\rho(P) = 1$ since $P\mathbf{1} = \mathbf{1}$, and the Perron–Frobenius theorem applied to $P^\top$ yields the existence and uniqueness (for irreducible P) of the stationary distribution π as the left Perron eigenvector, normalized to sum to 1. This is the linear-algebraic foundation of the existence and uniqueness results for stationary distributions established probabilistically in Chapter 3.

A.2 Diagonalization and Spectral Decomposition

Proposition: Spectral Decomposition of Reversible Chains

If P is reversible with respect to π, then P is self-adjoint with respect to the inner product $\langle f, g\rangle_\pi = \sum_x f(x)g(x)\pi(x)$, and therefore admits an orthogonal eigenbasis $\{\psi_k\}_{k=0}^{N-1}$ with real eigenvalues $1 = \lambda_0 > \lambda_1 \geq \cdots \geq \lambda_{N-1} \geq -1$.

This decomposition underlies every spectral bound on mixing times in Chapter 5: writing any initial distribution in the eigenbasis $\{\psi_k\}$, the action of P^n multiplies each component by λ_k^n, so all components with $|\lambda_k| < 1$ decay geometrically and the absolute spectral gap $\gamma^* = 1 - \max_{k\geq 1} |\lambda_k|$ controls the rate.

Remark: Jordan Form for Non-Reversible Chains

When P is not reversible (or not even diagonalizable), the Perron–Frobenius eigenvalue 1 remains simple for an irreducible aperiodic chain, but the remaining structure is governed by the Jordan canonical form rather than an orthogonal eigenbasis; mixing time bounds in this setting (Chapter 7) are typically obtained via coupling rather than direct spectral estimates, precisely because the spectral decomposition is less well-behaved.

A.3 Doubly Stochastic Matrices and the Birkhoff–von Neumann Theorem

A special and important class of stochastic matrices arises when both rows and columns sum to one; such matrices are exactly the transition matrices for which the uniform distribution is stationary, a fact used repeatedly in the random walk on graphs material of Chapter 11.

Definition: Doubly Stochastic Matrix

A matrix $P \in \mathbb{R}^{N\times N}$ is **doubly stochastic** if $P_{ij} \geq 0$ for all i, j, every row sums to 1, and every column sums to 1.

Proposition: Uniform Distribution is Stationary

If P is doubly stochastic, the uniform distribution $\pi(x) = 1/N$ is stationary for P.

Proof. $(\pi P)(y) = \sum_x \pi(x)P(x, y) = \frac{1}{N}\sum_x P(x, y) = \frac{1}{N}$, using the column-sum-one property in the last step, which gives $\pi P = \pi$. □

Theorem: Birkhoff–von Neumann

Every doubly stochastic matrix P is a convex combination of permutation matrices: there exist permutation matrices $\Pi_1, \dots, \Pi_k$ and weights $\lambda_1, \dots, \lambda_k \geq 0$ summing to 1 such that $P = \sum_{i=1}^{k} \lambda_i \Pi_i$. Moreover k can always be taken at most $N^2 - 2N + 2$.

Proof sketch. The set of doubly stochastic matrices forms a convex polytope (the **Birkhoff polytope**) inside $\mathbb{R}^{N \times N}$, cut out by the row-sum, column-sum, and nonnegativity constraints. The theorem asserts that the extreme points of this polytope are exactly the $N!$ permutation matrices. This is proved by showing that any doubly stochastic matrix that is not a permutation matrix has a positive entry off some permutation support, allowing a perturbation along a cycle in the bipartite graph of positive entries that decreases the number of positive entries while remaining inside the polytope; iterating this reduction (an instance of the same cycle-canceling idea used in network flow algorithms) terminates at a permutation matrix, and a convexity argument (Carathéodory's theorem applied to the polytope's dimension) bounds the number of terms needed in the convex combination. □

This decomposition gives a second, purely combinatorial proof that random walk on a doubly stochastic chain has uniform stationary distribution: viewing P as a probabilistic mixture of permutations, applying P once is equivalent to applying a uniformly chosen (according to the weights λ_i) permutation, and permutations trivially preserve the uniform distribution; the doubly stochastic property is exactly what allows the Perron–Frobenius eigenvector to be read off without any computation, in sharp contrast to general irreducible chains where finding π requires solving the full balance equations of Chapter 3.

Example: Random Walk on a Regular Graph is Doubly Stochastic

For simple random walk on a d-regular graph G, $P(x, y) = \mathbf{1}\{x \sim y\}/d$. Since every vertex has degree exactly d, each column of P also sums to $\sum_{x : x \sim y} 1/d = d/d = 1$, confirming double stochasticity directly from d-regularity, and recovering the well-known fact that the uniform distribution is stationary for random walk on any regular graph without needing to verify detailed balance explicitly.

Appendix B

Glossary of Notation and Standard Distributions

B.1 Notation Conventions

Throughout this volume:

- $\mathcal{X}$ denotes the (typically finite or countable) state space of a Markov chain.
- $P(x, y)$ or P_{xy} denotes the one-step transition probability from x to y; P^n denotes the n-step transition matrix.
- π denotes a stationary distribution; $\pi(x)$ its mass at state x.
- $\|\mu - \nu\|_{\mathrm{TV}}$ denotes total variation distance, $\|\mu - \nu\|_{\mathrm{TV}} = \frac{1}{2}\sum_x |\mu(x) - \nu(x)|$.
- $t_{\mathrm{mix}}(\varepsilon)$ denotes the ε-mixing time, t_{rel} the relaxation time $1/\gamma$.
- $\mathbb{E}_x[\cdot]$, $\mathbb{P}_x(\cdot)$ denote expectation and probability for the chain started at x.
- $T_A = \inf\{n \geq 0 : X_n \in A\}$ denotes the hitting time of a set A; T_A^+ the analogous first *return* time when $X_0 \in A$.
- γ (without subscript) denotes the spectral gap $1 - \lambda_1$; γ^* the absolute spectral gap $1 - \max(\lambda_1, |\lambda_{N-1}|)$.

B.2 Standard Discrete Distributions Used in Examples

Reference: Common Distributions

- **Bernoulli**(p): $\mathbb{P}(X = 1) = p$, $\mathbb{P}(X = 0) = 1 - p$, mean p, variance $p(1-p)$.
- **Binomial**(n, p): $\mathbb{P}(X = k) = \binom{n}{k} p^k (1-p)^{n-k}$, mean np, variance $np(1-p)$.
- **Geometric**(p) (on $\{1, 2, \ldots\}$): $\mathbb{P}(X = k) = (1-p)^{k-1}p$, mean $1/p$, variance $(1-p)/p^2$.

- **Poisson**(λ): $\mathbb{P}(X = k) = e^{-\lambda}\lambda^k/k!$, mean and variance both λ.
- **Exponential**(λ) (continuous): density $\lambda e^{-\lambda t}$, mean $1/\lambda$, the unique memoryless continuous distribution, fundamental to the construction of continuous-time chains in Chapters 9–10.

Appendix C

Pseudocode Reference for Key Algorithms

This appendix collects, in a uniform pseudocode format, the principal algorithms developed throughout the volume, for ease of implementation reference.

Algorithm: Power Iteration for the Stationary Distribution

Input: transition matrix P, tolerance ε.
Initialize: $\mu_0 \leftarrow$ uniform distribution on $\mathcal{X}$.
Repeat for $n = 0, 1, 2, \ldots$:
  $\mu_{n+1} \leftarrow \mu_n P$
  if $\|\mu_{n+1} - \mu_n\|_{\mathrm{TV}} < \varepsilon$: **return** μ_{n+1}
This is the direct algorithmic counterpart of the convergence theorem of Chapter 4, and its convergence rate is governed by the spectral gap exactly as in the bounds of Chapter 5.

Algorithm: Metropolis–Hastings (Reference Version)

Input: target distribution π (up to normalization), proposal kernel Q.
At state x**:**
  sample $y \sim Q(x, \cdot)$
  compute $\alpha(x, y) = \min\left(1, \dfrac{\pi(y)Q(y, x)}{\pi(x)Q(x, y)}\right)$
  with probability $\alpha(x, y)$: move to y; else: stay at x
See Chapter 10 for the proof that this construction yields a chain reversible with respect to π.

Algorithm: Coupling from the Past (CFTP)

Input: monotone update function ϕ, maximal and minimal states $\hat{0}, \hat{1}$.
For $T = 1, 2, 4, 8, \ldots$ (doubling):
  run coupled chains from $\hat{0}$ and $\hat{1}$ over time window $[-T, 0]$ using a fixed random map sequence $(\phi_{-T}, \ldots, \phi_{-1})$

if both chains coalesce by time 0: **return** the common value

See Chapter 21 (and the Propp–Wilson references in the Bibliography) for the proof of exactness and the monotonicity condition required for the two-chain reduction.

Algorithm: Forward Algorithm for HMM Filtering (Reference Version)

Input: transition matrix P, emission kernel g, observations $y_0, \dots, y_n$.
Initialize: $\alpha_0(x) \propto \mu_0(x) g(y_0 \mid x)$.
For $k = 0, \dots, n-1$:
$\alpha_{k+1}(x) \propto g(y_{k+1} \mid x) \sum_{x'} P(x', x)\, \alpha_k(x')$
Return α_n, proportional to the filtering distribution π_n.

See Chapter 22 for the derivation and the related backward and Viterbi recursions.

Appendix D

Compendium of Mixing Time Bounds and Concentration Inequalities

D.1 Master Table of Mixing Time Bounds

This appendix gathers, side by side, the principal upper and lower bound techniques for mixing times developed across Chapters 5 through 21, to serve as a quick-reference summary and to make explicit how the different approaches relate to one another.

Summary: Spectral Bounds (Chapter 5)

For reversible P with stationary distribution π and spectral gap $\gamma = 1 - \lambda_1$:

$$t_{\text{mix}}(\varepsilon) \;\le\; \frac{1}{\gamma}\log\!\left(\frac{1}{\varepsilon\pi_{\min}}\right), \qquad t_{\text{rel}} = \frac{1}{\gamma^*} \;\le\; t_{\text{mix}}(\varepsilon)\cdot O(\log(1/\varepsilon))^{-1}.$$

Sharpest when the full spectrum is known or well-approximated; requires reversibility for the cleanest statement, though analogous singular-value bounds exist for non-reversible chains.

Summary: Coupling Bounds (Chapters 6–7)

For a coupling (X_n, Y_n) with coalescence time τ_{couple}:

$$\|P^n(x,\cdot) - P^n(y,\cdot)\|_{\text{TV}} \le \mathbb{P}(\tau_{\text{couple}} > n), \qquad t_{\text{mix}}(\varepsilon) \le \min\{n : \max_{x,y}\mathbb{P}(\tau_{\text{couple}} > n) \le \varepsilon\}.$$

Most effective when an explicit, analyzable coupling exists (e.g. monotone couplings, path coupling); often gives sharp results even when the spectral gap is hard to compute directly.

Summary: Conductance / Cheeger Bounds (Chapter 5)

With conductance $\Phi = \min_{S:\pi(S)\le 1/2} \frac{Q(S,S^c)}{\pi(S)}$:

$$\frac{\Phi^2}{2} \;\le\; \gamma \;\le\; 2\Phi.$$

Useful for proving *slow* mixing (via a small-conductance bottleneck set) as well as fast mixing (via isoperimetric/expansion arguments), and is the bridge between the purely algebraic spectral gap and the geometric structure of the state space.

Summary: Log-Sobolev and Functional Inequality Bounds

With log-Sobolev constant α:

$$t_{\text{mix}}(\varepsilon) \;\le\; \frac{1}{2\alpha}\log\log\left(\frac{1}{\pi_{\min}}\right) + \frac{1}{2\gamma}\log\left(\frac{2}{\varepsilon^2}\right).$$

Captures the distinction between the spectral-gap-only bound (which has a $\log(1/\pi_{\min})$ term) and sharper bounds that replace one log factor with a double logarithm, explaining cutoff phenomena (Chapter 14) in many high-dimensional examples (e.g. the random transposition and Glauber dynamics chains).

D.2 Concentration Inequalities Used Throughout the Volume

Reference: Markov and Chebyshev Inequalities

For nonnegative X and $a > 0$: $\mathbb{P}(X \ge a) \le \mathbb{E}[X]/a$ (Markov). For X with finite variance: $\mathbb{P}(|X - \mathbb{E}X| \ge a) \le \mathrm{Var}(X)/a^2$ (Chebyshev). These elementary bounds underlie the first-moment recurrence/transience criteria of Chapter 3.

Reference: Chernoff Bound

For $X = \sum_{i=1}^{n} X_i$ a sum of independent Bernoulli(p_i) random variables with mean $\mu = \mathbb{E}[X]$:

$$\mathbb{P}(X \ge (1+\delta)\mu) \le \exp\left(-\frac{\delta^2\mu}{2+\delta}\right), \qquad \delta > 0.$$

This is the prototype large-deviation bound (Chapter 20), recovered as the $i.i.d.$ special case of the Markov chain large deviation principle when transitions are independent of the current state.

Reference: Azuma–Hoeffding Inequality

For a martingale (M_n) with bounded increments $|M_k - M_{k-1}| \le c_k$:

$$\mathbb{P}(|M_n - M_0| \ge t) \le 2\exp\left(-\frac{t^2}{2\sum_{k=1}^{n} c_k^2}\right).$$

Used throughout Chapter 16 to control fluctuations of martingale functionals of Markov chains, and a key technical tool in proving concentration of the mixing time itself around its expectation in some chain families.

Reference: McDiarmid's Bounded Differences Inequality

If $f(x_1, \dots, x_n)$ changes by at most c_i when the i-th coordinate is altered (other coordinates fixed), and $X_1, \dots, X_n$ are independent:

$$\mathbb{P}(|f(X_1, \dots, X_n) - \mathbb{E}f| \ge t) \le 2\exp\left(-\frac{2t^2}{\sum_i c_i^2}\right).$$

A generalization of Azuma–Hoeffding via the Doob martingale associated with f; useful for concentration of statistics computed from many weakly dependent samples of a rapidly mixing chain, after an appropriate mixing-time correction to the effective sample size.

Remark: Concentration Under Mixing

For a stationary, ρ-mixing sequence (a chain with spectral gap γ), the effective number of independent samples in n observations scales as $n\gamma$ rather than n, and the concentration inequalities above can be adapted by replacing n with $n\gamma$ in the denominator of the exponent, formalizing the heuristic that mixing time bounds and concentration inequalities are two sides of the same coin: both quantify how quickly a Markov chain's trajectory statistics resemble those of an i.i.d. sample from π.

D.3 Cross-Reference Table: Where Each Technique Is Developed

Cross-Reference

- Spectral methods and Cheeger's inequality: Chapter 5, with advanced topics (expander mixing, Ramanujan graphs, KLS) in Chapter 18.
- Coupling and path coupling: Chapters 6–7, with monotone coupling and CFTP in Chapter 19.
- Functional inequalities (Poincaré, log-Sobolev): referenced in Chap-

ter 14's cutoff discussion and Appendix D above.

- Large deviations and rare-event asymptotics: Chapter 20.
- Renewal-theoretic methods: Chapter 19.
- Electrical network / resistance methods: Chapter 21.
- Martingale and stopping time methods: Chapter 16, with optimal stopping MDPs in Chapter 23.

This table is intended as a navigational aid for readers returning to the volume to look up a specific technique for use in their own research or coursework.

Appendix E

Selected Exercises

This closing appendix collects a curated set of exercises spanning the volume, organized by part, intended for self-study or course use. Solutions are not provided; each exercise is designed to be solvable using only the theory developed up to and including the chapter referenced.

E.1 Foundations and Basic Structure (Chapters 1–3)

Exercise 1

Let P be the transition matrix of a random walk on the cycle $\mathbb{Z}/n\mathbb{Z}$ that moves $+1$ or -1 with probability $1/2$ each. Show directly from the definition that the uniform distribution is stationary, and determine the period of the chain as a function of n.

Exercise 2

A Markov chain on $\{1, 2, 3, 4, 5\}$ has $P(1, 2) = P(2, 3) = P(3, 4) = P(4, 5) = 1$ and $P(5, 1) = 1$. Classify all states as recurrent or transient, determine the period, and find the unique stationary distribution.

Exercise 3

Using the first-step decomposition technique of Chapter 3, compute the expected hitting time $\mathbb{E}_0[T_n]$ for simple random walk on $\{0, 1, \ldots, n\}$ with absorbing/reflecting boundary conditions of your choice, and verify your formula against the $n = 2$ and $n = 3$ cases by direct enumeration.

E.2 Stationarity, Ergodicity, and Spectral Theory (Chapters 3–5)

Exercise 4

Show that for a reversible chain, all eigenvalues of P are real, using the symmetrization argument of Chapter 5 (conjugating P by $\mathrm{diag}(\sqrt{\pi(x)})$ into a symmetric matrix).

Exercise 5

Compute the full spectrum of the transition matrix for simple random walk on the complete graph K_n (move to a uniformly random other vertex), and use it to give an exact (not just asymptotic) formula for $\|P^t(x,\cdot) - \pi\|_{\mathrm{TV}}$.

Exercise 6

Prove Cheeger's inequality $\gamma \le 2\Phi$ (the easy direction) by exhibiting a test function in the variational characterization of γ built from the indicator of the minimizing conductance set.

E.3 Coupling and Mixing Time Bounds (Chapters 6–7, 14)

Exercise 7

Construct an explicit coupling for the lazy random walk on the hypercube $\{0,1\}^n$ (at each step, pick a uniformly random coordinate and resample it uniformly) and use it to derive an $O(n \log n)$ mixing time bound.

Exercise 8

For the chain in Exercise 7, verify numerically (for small n, e.g. $n = 4$ or $n = 6$) that the mixing time exhibits a cutoff at $\frac{1}{2} n \log n$ with window $O(n)$, by computing $\|P^t(\mathbf{0},\cdot) - \pi\|_{\mathrm{TV}}$ at several values of t.

Exercise 9

Use path coupling to bound the mixing time of the Glauber dynamics for proper colorings of a graph with maximum degree Δ, when the number of colors q satisfies $q > 2\Delta$, and identify exactly where the condition $q > 2\Delta$ enters the contraction argument.

E.4 Continuous-Time Chains and MCMC (Chapters 8–10)

Exercise 10

Derive the generator matrix Q for a continuous-time birth-death chain with birth rate $\lambda_n = \lambda$ and death rate $\mu_n = \mu n$ (an $M/M/\infty$ queue), and find its stationary distribution by solving $\pi Q = 0$.

Exercise 11

Implement (on paper or in code) the Metropolis–Hastings algorithm to sample from a discretized Gaussian target on $\{-N, \dots, N\}$ using a simple random walk proposal, and empirically estimate the mixing time for $N = 20$ by monitoring the autocorrelation of a test statistic.

E.5 Martingales, Statistical Physics, and General State Spaces (Chapters 16–17)

Exercise 12

Use the optional stopping theorem to compute the gambler's ruin probability for a biased random walk on $\{0, 1, \dots, N\}$ with absorbing boundaries, by finding an appropriate martingale (a function of the walk that is itself a martingale).

Exercise 13

For the Ising model on the cycle $\mathbb{Z}/n\mathbb{Z}$ at inverse temperature β, compute the partition function exactly using the transfer matrix method, and relate the transfer matrix's spectral gap to the correlation length of the model.

E.6 Advanced Topics: Renewal, Large Deviations, Random Graphs, HMMs, and MDPs (Chapters 19–23)

Exercise 14

Verify the renewal equation $u_n = \sum_{k=1}^{n} f_k u_{n-k}$ directly for the return-time distribution of simple random walk on $\mathbb{Z}$ reflected at 0 and N, for small N, by explicit enumeration of paths.

Exercise 15

For the two-state chain of Chapter 20's worked example, compute the rate function $I(x)$ explicitly (not just its leading-order form) and verify that I attains its minimum value 0 exactly at $x = \pi(1)$.

Exercise 16

Using the electrical network correspondence of Chapter 21, compute the effective resistance between two adjacent vertices of the $n \times n$ grid graph and use it to determine whether simple random walk on the infinite grid $\mathbb{Z}^2$ is recurrent, recovering Pólya's theorem for $d = 2$ by this alternative route.

Exercise 17

For a hidden Markov model with a two-state hidden chain and Gaussian emissions, derive the forward algorithm recursion explicitly in terms of Gaussian densities, and explain why, unlike the finite-emission case, the resulting filtering distribution remains expressible in closed form only when the hidden state space itself is finite (contrast with the Kalman filter, which handles continuous hidden states under linear-Gaussian assumptions).

Exercise 18

Formulate the classical "secretary problem" (observe n candidates in random order, must accept or reject immediately, want to maximize the probability of selecting the best) as an optimal stopping MDP in the style of Chapter 23, and rederive the well-known $1/e$ asymptotic success probability using the superharmonic majorant characterization of the value function.

Exercise 19

Let (X_n) be an irreducible Markov chain on a finite state space with stationary distribution π and let f be a function with $\mathbb{E}_\pi[f] = 0$. Using the renewal-reward theorem of Chapter 19 together with the central limit theorem for Markov chains, give two independent derivations of the asymptotic variance formula $\sigma_f^2 = \mathrm{Var}_\pi(f) + 2\sum_{k=1}^{\infty} \mathrm{Cov}_\pi(f(X_0), f(X_k))$, one via the regenerative structure at returns to a fixed state, and one via the spectral decomposition of P when the chain is reversible. Compare the two derivations and identify which hypotheses each requires.

Exercise 20

Consider a queueing network modeled as a continuous-time Markov chain with two tandem $M/M/1$ queues (Chapter 13), where the departure process of the first queue is the arrival process of the second. Show that, in equilibrium, the departure process from the first queue is itself a Poisson process (Burke's theorem), and use this fact together with the product-form stationary distribution

to compute the mean total number of customers in the network as a function of the arrival rate λ and the two service rates μ_1, μ_2.

Remark: On the Use of These Exercises

The exercises above are graded roughly by difficulty within each section, and several (notably Exercises 8, 11, and the numerical verification in Exercise 16) are best approached computationally; readers are encouraged to implement small simulations to build intuition before attempting the corresponding analytical derivations, in keeping with the volume's recurring emphasis on the interplay between exact theory and concrete, checkable examples.

Bibliography

[1] D. Aldous and J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, 2002.

[2] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*, 2nd ed., AMS, 2017.

[3] P. Diaconis and L. Saloff-Coste, *Logarithmic Sobolev inequalities for finite Markov chains*, Ann. Appl. Probab. 6 (1996), 695–750.

[4] J. S. Rosenthal, *Minorization Conditions and Convergence Rates for MCMC*, JASA 90 (1995), 558–566.

[5] R. Montenegro and P. Tetali, *Mathematical Aspects of Mixing Times in Markov Chains*, Found. Trends TCS 1 (2006), 237–354.

[6] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica 8 (1988), 261–277.

[7] J. Propp and D. Wilson, *Exact sampling with coupled Markov chains*, Random Struct. Algorithms 9 (1996), 223–252.

[8] S. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*, 2nd ed., CUP, 2009.

[9] P. Diaconis and M. Shahshahani, *Generating a random permutation with random transpositions*, Z. Wahrsch. 57 (1981), 159–179.

[10] M. Jerrum and A. Sinclair, *Approximating the permanent*, SIAM J. Comput. 18 (1989), 1149–1178.

[11] S. Asmussen, *Applied Probability and Queues*, 2nd ed., Springer, 2003.

[12] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed., Springer, 1998.

[13] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, 3rd ed., Springer, 2012.

[14] R. Lyons and Y. Peres, *Probability on Trees and Networks*, CUP, 2016.

[15] P. G. Doyle and J. L. Snell, *Random Walks and Electric Networks*, MAA, 1984.

[16] L. R. Rabiner, *A tutorial on hidden Markov models and selected applications in speech recognition*, Proc. IEEE 77 (1989), 257–286.

[17] A. Doucet, N. de Freitas, and N. Gordon (eds.), *Sequential Monte Carlo Methods in Practice*, Springer, 2001.

[18] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley, 1994.

[19] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, 2018.